

Specification: The Pattern That Signifies Intelligence

By William A. Dembski

August 15, 2005, version 1.22

1	Specification as a Form of Warrant	1
2	Fisherian Significance Testing	2
3	Specifications via Probability Densities	5
4	Specifications via Compressibility	9
5	Prespecifications vs. Specifications	12
6	Specificity	15
7	Specified Complexity	20
8	Design Detection	25
	Acknowledgment	31
	Addendum 1: Note to Readers of <i>TDI & NFL</i>	32
	Addendum 2: Bayesian Methods	35
	Endnotes	38

ABSTRACT: Specification denotes the type of pattern that highly improbable events must exhibit before one is entitled to attribute them to intelligence. This paper analyzes the concept of specification and shows how it applies to design detection (i.e., the detection of intelligence on the basis of circumstantial evidence). Always in the background throughout this discussion is the fundamental question of Intelligent Design (ID): Can objects, even if nothing is known about how they arose, exhibit features that reliably signal the action of an intelligent cause? This paper reviews, clarifies, and extends previous work on specification in my books *The Design Inference* and *No Free Lunch*.

1 Specification as a Form of Warrant

In the 1990s, Alvin Plantinga focused much of his work on the concept of *warrant*. During that time, he published three substantial volumes on that topic.¹ For Plantinga, warrant is what turns true belief into knowledge. To see what is at stake with warrant, consider the following remark by Bertrand Russell on the nature of error:

Error is not only the absolute error of believing what is false, but also the quantitative error of believing more or less strongly than is warranted by the degree of credibility properly attaching to the proposition believed in relation to the believer's knowledge. A man who is quite convinced that a certain horse will win the Derby is in error even if the horse does win.²

It is not enough merely to believe that something is true; there also has to be something backing up that belief. Accordingly, belief and persuasion are intimately related concepts: we believe that

about which we are persuaded. This fact is reflected in languages for which the same word denotes both belief and persuasion. For instance, in ancient Greek, the verb *peitho* denoted both to persuade and to believe.³

Over the last decade, much of my work has focused on design detection, that is, sifting the effects of intelligence from material causes. Within the method of design detection that I have developed, *specification* functions like Plantinga's notion of warrant: just as for Plantinga warrant is what must be added to true belief before one is entitled to call it knowledge, so within my framework specification is what must be added to highly improbable events before one is entitled to attribute them to design. The connection between specification and warrant is more than a loose analogy: specification constitutes a probabilistic form of warrant, transforming the suspicion of design into a warranted belief in design.

My method of design detection, first laid out in *The Design Inference* and then extended in *No Free Lunch*,⁴ provides a systematic procedure for sorting through circumstantial evidence for design. This method properly applies to circumstantial evidence, not to smoking-gun evidence. Indeed, such methods are redundant in the case of a smoking gun where a designer is caught red-handed. To see the difference, consider the following two cases:

- (1) Your daughter pulls out her markers and a sheet of paper at the kitchen table while you are eating breakfast. As you're eating, you see her use those markers to draw a picture. When you've finished eating, she hands you the picture and says, "This is for you, Daddy."
- (2) You are walking outside and find an unusual chunk of rock. You suspect it might be an arrowhead. Is it truly an arrowhead, and thus the result of design, or is it just a random chunk of rock, and thus the result of material forces that, for all you can tell, were not intelligently guided?

In the first instance, the design and designer couldn't be clearer. In the second, we are dealing with purely circumstantial evidence: we have no direct knowledge of any putative designer and no direct knowledge of how such a designer, if actual, fashioned the item in question. All we see is the pattern exhibited by the item. Is it the sort of pattern that reliably points us to an intelligent source? That is the question my method of design detection asks and that the concept of specification answers.

2 Fisherian Significance Testing

A variant of specification, sometimes called *prespecification*, has been known for at least a century. In the late 1800s, the philosopher Charles Peirce referred to such specifications as *predesignations*.⁵ Such specifications are also the essential ingredient in Ronald Fisher's theory of statistical significance testing, which he devised in the first half of the twentieth century. In

Fisher's approach to testing the statistical significance of hypotheses, one is justified in rejecting (or eliminating) a chance hypothesis provided that a sample falls within a prespecified *rejection region* (also known as a *critical region*).⁶ For example, suppose one's chance hypothesis is that a coin is fair. To test whether the coin is biased in favor of heads, and thus not fair, one can set a rejection region of ten heads in a row and then flip the coin ten times. In Fisher's approach, if the coin lands ten heads in a row, then one is justified rejecting the chance hypothesis.

Fisher's approach to hypothesis testing is the one most widely used in the applied statistics literature and the first one taught in introductory statistics courses. Nevertheless, in its original formulation, Fisher's approach is problematic: for a rejection region to warrant rejecting a chance hypothesis, the rejection region must have sufficiently small probability. But how small is small enough? Given a chance hypothesis and a rejection region, how small does the probability of the rejection region have to be so that if a sample falls within it, then the chance hypothesis can legitimately be rejected? Fisher never answered this question. The problem here is to justify what is called a *significance level* such that whenever the sample falls within the rejection region and the probability of the rejection region given the chance hypothesis is less than the significance level, then the chance hypothesis can be legitimately rejected.

More formally, the problem is to justify a significance level α (always a positive real number less than one) such that whenever the sample (an event we will call E) falls within the rejection region (call it T) and the probability of the rejection region given the chance hypothesis (call it \mathbf{H}) is less than α (i.e., $\mathbf{P}(T|\mathbf{H}) < \alpha$), then the chance hypothesis \mathbf{H} can be rejected as the explanation of the sample. In the applied statistics literature, it is common to see significance levels of .05 and .01. The problem to date has been that any such proposed significance levels have seemed arbitrary, lacking "a rational foundation."⁷

In *The Design Inference*, I show that significance levels cannot be set in isolation but must always be set in relation to the probabilistic resources relevant to an event's occurrence.⁸ In the context of Fisherian significance testing, probabilistic resources refer to the number opportunities for an event to occur. The more opportunities for an event to occur, the more possibilities for it to land in the rejection region and thus the greater the likelihood that the chance hypothesis in question will be rejected. It follows that a seemingly improbable event can become quite probable once enough probabilistic resources are factored in. Yet if the event is sufficiently improbable, it will remain improbable even after all the available probabilistic resources have been factored in.

Critics of Fisher's approach to hypothesis testing are therefore correct in claiming that significance levels of .05, .01, and the like that regularly appear in the applied statistics literature are arbitrary. A significance level is supposed to provide an upper bound on the probability of a rejection region and thereby ensure that the rejection region is sufficiently small to justify the elimination of a chance hypothesis (essentially, the idea is to make a target so small that an archer is highly unlikely to hit it by chance). Rejection regions eliminate chance hypotheses when events that supposedly happened in accord with those hypotheses fall within the rejection

regions. The problem with significance levels like .05 and .01 is that typically they are instituted without reference to the probabilistic resources relevant to controlling for false positives.⁹ Once those probabilistic resources are factored in, however, setting appropriate significance levels becomes straightforward.¹⁰

To illustrate the use of probabilistic resources here, imagine the following revision of the criminal justice system. From now on, a convicted criminal serves time in prison until he flips k heads in a row, where this number is selected according to the severity of the offense (the worse the offense, the bigger k). Let's assume that all coin flips are fair and duly recorded — no cheating is possible. Thus, for a 10-year prison sentence, if we assume the prisoner can flip a coin once every five seconds (this seems reasonable), the prisoner will perform 12 tosses per minute or 720 per hour or 5,760 in an eight-hour work day or 34,560 in a six-day work week or 1,797,120 in a year or (approximately) 18 million in ten years. Of these, half will, on average, be heads. Of these, in turn, half will, on average, be followed by heads. Continuing in this vein we find that a sequence of 18 million coin tosses yields 23 heads in a row roughly half the time and strictly less than 23 heads in a row the other half (note that 2^{23} is about 9 million, which is about half of the total number of coin tosses). Thus, if we required a prisoner to flip 23 heads in a row before being released, we would, on average, expect to see him freed within approximately 10 years. Of course, specific instances will vary: some prisoners will be released after only a short stay, whereas others never will record the elusive 23 heads in a row.

Flipping 23 heads in a row is unpleasant but doable. On the other hand, for a prisoner to flip 100 heads in a row with a fair coin is effectively impossible. The probability of getting 100 heads in a row on a given trial is so small that the prisoner has no practical hope of getting out of prison. And this remains true even if his life expectancy and coin-tossing ability were substantially increased. If he could, for instance, make 10^{10} (i.e., 10 billion) attempts each year to obtain 100 heads in a row (this is coin-flipping at a rate of over 500 coin flips per second every hour of every day for a full year), then he stands only an even chance of getting out of prison in 10^{20} years (i.e., a hundred billion billion years). His probabilistic resources are so inadequate for obtaining the desired 100 heads that it is pointless for him to entertain hopes of freedom.

Requiring a prisoner to toss 23 heads in a row before being released from prison corresponds, in Fisherian terms, to a significance level of roughly 1 in 9 million (i.e., $\alpha = 1/9,000,000$). Accordingly, before we attribute irregularity to the prisoner's coin tossing (e.g., by claiming that the coin is biased or that there was cheating), we need to see more than 23 heads in a row. How many more? That depends on the degree of warrant we require to confidently eliminate chance. All chance elimination arguments are fallible in the sense that an event might land in a low probability rejection region even if the chance hypothesis holds. But at some point the improbability becomes palpable and we look to other hypotheses. With 100 heads in a row, that is obviously the case: the probability of getting 100 heads in a row is roughly 1 in 10^{30} , which is drastically smaller than 1 in 9 million. Within Fisher's theory of statistical significance testing, a prespecified event of such small probability is enough to disconfirm the chance hypothesis.

3 Specifications via Probability Densities

In the logic of Fisher’s approach to statistical significance testing, for the chance occurrence of an event E to be legitimately rejected because E falls in a rejection region T , T had to be identified *prior to* the occurrence of E . This is to avoid the familiar problem known among statisticians as “data snooping” or “cherry picking,” in which a pattern (in this case T) is imposed on an event (in this case the sample E) after the fact. Requiring the rejection region to be set prior to the occurrence of E safeguards against attributing patterns to E that are factitious and that do not properly preclude the occurrence of E by chance. This safeguard, however, is unduly restrictive. Indeed, even within Fisher’s approach to eliminating chance hypotheses, this safeguard is relaxed.

To see this, consider that a reference class of possibilities Ω , for which patterns and events can be defined, usually comes with some additional geometric structure together with a privileged (probability) measure that preserves that geometric structure. In case Ω is finite or countably infinite, this measure is typically just the counting measure (i.e., it counts the number of elements in a given subset of Ω ; note that normalizing this measure on finite sets yields a uniform probability). In case Ω is uncountable but suitably bounded (i.e., compact), this privileged measure is a uniform probability. In case Ω is uncountable and unbounded, this privileged measure becomes a uniform probability when restricted to and normalized with respect to the suitably bounded subsets of Ω . Let us refer to this privileged measure as \mathbf{U} .¹¹

Now the interesting thing about \mathbf{U} in reference to Fisher’s approach to eliminating chance hypotheses is that it allows probabilities of the form $\mathbf{P}(\cdot|\mathbf{H})$ (i.e., those defined explicitly with respect to a chance hypothesis \mathbf{H}) to be represented as probabilities of the form $f \cdot d\mathbf{U}$ (i.e., the product of a nonnegative function f , known as a *probability density function*, and the measure \mathbf{U}). The “ d ” in front of \mathbf{U} here signifies that to evaluate this probability requires integrating f with respect to \mathbf{U} .¹² The technical details here are not important. What is important is that within Fisher’s approach to hypothesis testing the probability density function f is used to identify rejection regions that in turn are used to eliminate chance. These rejection regions take the form of *extremal sets*, which we can represent as follows (γ and δ are real numbers):

$$T_\gamma = \{\omega \in \Omega \mid f(\omega) \geq \gamma\},$$

$$T_\delta = \{\omega \in \Omega \mid f(\omega) \leq \delta\}.$$


T_γ consists of all possibilities in the reference class Ω for which the density function f is at least γ . Similarly, T_δ consists of all possibilities in the reference class Ω for which the density function f is no more than δ . Although this way of characterizing rejection regions may seem abstract, what is actually going on here is simple. Any probability density function f is a nonnegative real-valued function defined on Ω . As such, f induces what might be called a *probability landscape* (think of Ω as a giant plane and f as describing the elevation of a landscape over Ω). Where f is high corresponds to where the probability measure $f \cdot d\mathbf{U}$ concentrates a lot of probability. Where f

is low corresponds to where the probability measure $f dU$ concentrates little probability. Since f cannot fall below zero, we can think of the landscape as never dipping below sea-level. T_δ then corresponds to those places in Ω where the probability landscape induced by f is no higher than δ whereas T_γ corresponds to those places in Ω where probability landscape is at least as high as γ . For a bell-shaped curve, T_γ corresponds to the region under the curve where it approaches a maximum, and T_δ corresponds to the tails of the distribution.

Since this way of characterizing rejection regions may still seem overly abstract, let us apply it to several examples. Suppose, for instance, that Ω is the real line and that the hypothesis \mathbf{H} describes a normal distribution with, let us say, mean zero and variance one. In that case, for the probability measure $f dU$ that corresponds to the $\mathbf{P}(\cdot|\mathbf{H})$, the density function f has the form

$$f(x) = \frac{1}{\sqrt{2\pi}} e^{-\frac{x^2}{2}}.$$

This function is everywhere positive and attains its maximum at the mean (i.e., at $x = 0$, where f takes the value $1/\sqrt{2\pi}$, or approximately .399). At two standard deviations from the mean (i.e., for the absolute value of x at least 2), this function attains but does not exceed .054. Thus for $\delta = .054$, T_δ corresponds to two tails of the normal probability distribution:



Moreover, since those tails are at least two standard deviations from the mean, it follows that $\mathbf{P}(T_\delta|\mathbf{H}) < .05$. Thus, for a significance level $\alpha = .05$ and a sample E that falls within T_δ , Fisher's approach rejects attributing E to the chance hypothesis \mathbf{H} .

The important thing to note in this example is not the precise level of significance that was set (here $\alpha = .05$) or the precise form of the probability distribution under consideration (here a normal distribution with mean zero and variance one). These were simply given for concreteness. Rather, the important thing here is that the temporal ordering of rejection region and sample, where a rejection region (here T_δ) is first specified and a sample (here E) is then taken, simply was not an issue. For a statistician to eliminate the chance hypothesis \mathbf{H} as an explanation of E , it is not necessary for the statistician first to identify T_δ explicitly, then perform

an experiment to elicit E , and finally determine whether E falls within T_δ . \mathbf{H} , by inducing a probability density function f , automatically also inducing the extremal set T_δ which, given the significance level $\alpha = .05$, is adequate for eliminating \mathbf{H} once a sample E falls within T_δ . If you will, the rejection region T_δ comes automatically with the chance hypothesis \mathbf{H} , making it unnecessary to identify it explicitly prior to the sample E . Eliminating chance when samples fall in such extremal sets is standard statistical practice. Thus, in the case at hand, if a sample falls sufficiently far out in the tails of a normal distribution, the distribution is rejected as inadequate to account for the sample.

In this last example we considered extremal sets of the form T_δ at which the probability density function concentrates minimal probability. Nonetheless, extremal sets of the form T' at which the probability density function concentrates maximal probability can also serve as rejection regions within Fisher's approach, and this holds even if such regions are not explicitly identified prior to an experiment. Consider, for instance, the following example. Imagine that a die is to be thrown 6,000,000 times. Within Fisher's approach, the "null" hypothesis (i.e., the chance hypothesis most naturally associated with this probabilistic set-up) is a chance hypothesis \mathbf{H} according to which the die is fair (i.e., each face has probability $1/6$) and the die rolls are stochastically independent (i.e., one roll does not affect the others). Consider now the reference class of possibilities Ω consisting of all 6-tuples of nonnegative integers that sum to 6,000,000. The 6-tuple (6,000,000, 0, 0, 0, 0, 0) would, for instance, correspond to tossing the die 6,000,000 times with the die each time landing "1." The 6-tuple (1,000,000, 1,000,000, 1,000,000, 1,000,000, 1,000,000, 1,000,000), on the other hand, would correspond to tossing the die 6,000,000 times with each face landing exactly as often as the others. Both these 6-tuples represent possible outcomes for tossing the die 6,000,000 times, and both belong to Ω .

Suppose now that the die is thrown 6,000,000 times and that each face appears *exactly* 1,000,000 times (as in the previous 6-tuple). Even if the die is fair, there is something anomalous about getting *exactly* one million appearances of each face of the die. What's more, Fisher's approach is able to make sense of this anomaly. The probability distribution that \mathbf{H} induces on Ω is known as the multinomial distribution,¹³ and for each 6-tuple $(x_1, x_2, x_3, x_4, x_5, x_6)$ in Ω , it assigns a probability of

$$f(x_1, x_2, x_3, x_4, x_5, x_6) = \frac{6,000,000!}{x_1!x_2!x_3!x_4!x_5!x_6!} (1/6)^{6,000,000}$$

Although the combinatorics involved with the multinomial distribution are complicated (hence the common practice of approximating it with continuous probability distributions like the chi-square distribution), the reference class of possibilities Ω , though large, is finite, and the cardinality of Ω (i.e., the number of elements in Ω), denoted by $|\Omega|$, is well-defined (its order of magnitude is around 10^{33}).

The function f defines not only a probability for individual elements $(x_1, x_2, x_3, x_4, x_5, x_6)$ of Ω but also a probability density with respect to the counting measure \mathbf{U} (i.e., $\mathbf{P}(\cdot|\mathbf{H}) = f d\mathbf{U}$), and attains its maximum at the 6-tuple for which each face of the die appears exactly 1,000,000 times. Let us now take γ to be the value of $f(x_1, x_2, x_3, x_4, x_5, x_6)$ at $x_1 = x_2 = x_3 = x_4 = x_5 = x_6 = 1,000,000$ (γ is approximately 2.475×10^{-17}). Then T^γ defines a rejection region corresponding to those elements of Ω where f attains at least the value γ . Since γ is the maximum value that f attains and since there is only one 6-tuple where this value is attained (i.e., at $x_1 = x_2 = x_3 = x_4 = x_5 = x_6 = 1,000,000$), it follows that T^γ contains only one member of Ω and that the probability of T^γ is $\mathbf{P}(T^\gamma|\mathbf{H}) = f(1,000,000, 1,000,000, 1,000,000, 1,000,000, 1,000,000, 1,000,000)$, which is approximately 2.475×10^{-17} . This rejection region is therefore highly improbable, and its improbability will in most practical applications be far more extreme than any significance level we happen to set. Thus, if we observed exactly 1,000,000 appearances of each face of the die, we would, according to Fisher's approach, conclude that the die was not thrown by chance. Unlike the previous example, in which by falling in the tails of the normal distribution a sample deviated from expectation by too much, the problem here is that the sample matches expectation too closely. This is a general feature of Fisherian significance testing: samples that fall in rejection regions of the form T_δ deviate from expectation too much whereas samples that fall in rejection regions of the form T^γ match expectation too closely.

The problem of matching expectation too closely is as easily handled within Fisher's approach as the problem of deviating from expectation too much. Fisher himself recognized as much and saw rejection regions of the form T^γ as useful for uncovering data falsification. For instance, Fisher uncovered a classic case of data falsification in analyzing Gregor Mendel's data on peas. Fisher inferred that "Mendel's data were massaged," as one statistics text puts it, because Mendel's data matched his theory too closely.¹⁴ The match that elicited this charge of data falsification was a specified event whose probability was no more extreme than 1 in a 100,000 (a probability that is huge compared to the 1 in 10^{17} probability computed in the preceding example). Fisher concluded by charging Mendel's gardening assistant with deception.

To sum up, given the chance hypothesis \mathbf{H} , the extremal sets T^γ and T_δ were defined with respect to the probability density function f where $\mathbf{P}(\cdot|\mathbf{H}) = f d\mathbf{U}$. Moreover, given a significance level α and a sample E that falls within either of these extremal sets, Fisher's approach eliminates \mathbf{H} provided that the extremal sets have probability less than α . An obvious question now arises: In using either of the extremal sets T^γ or T_δ to eliminate the chance hypothesis \mathbf{H} , what is so special about basing these extremal sets on the probability density function f associated with the probability measure $\mathbf{P}(\cdot|\mathbf{H}) (= f d\mathbf{U})$? Why not let f be an arbitrary real-valued function?

Clearly, some restrictions must apply to f . If f can be arbitrary, then for any subset A of the reference class Ω , we can define f to be the indicator function 1_A (by definition this function equals 1 whenever an element of Ω is in A and 0 otherwise).¹⁵ For such an f , if $\gamma = 1$, then $T^\gamma = \{\omega \in \Omega \mid f(\omega) \geq \gamma\} = A$, and thus any subset of Ω becomes an extremal set for some function f . But this would allow Fisher's approach to eliminate any chance hypothesis whatsoever: For any sample E (of sufficiently small probability), find a subset A of Ω that includes E and such that

$\mathbf{P}(A|\mathbf{H})$ is less than whatever significance level α was decided. Then for $f = 1_A$ and $\gamma = 1$, $T^\gamma = A$. Consequently, if Fisher's approach extends to arbitrary f , \mathbf{H} would always have to be rejected. But by thus eliminating all chance hypotheses, Fisher's approach becomes useless.

What restrictions on f will therefore safeguard it from frivolously disqualifying chance hypotheses? Recall that initially it was enough for rejection regions to be given in advance of an experiment. Then we noted that the rejection regions corresponding to extremal sets associated with the probability density function associated with the chance hypothesis in question also worked (regardless of when they were identified — in particular, they did not have to be identified in advance of an experiment). Finally, we saw that placing no restrictions on the function f eliminated chance willy-nilly and thus could not legitimately be used to eliminate chance. But why was the complete lack of restrictions inadequate for eliminating chance? The problem — and this is why statisticians are so obsessive about making sure hypotheses are formulated in advance of experiments — is that with no restriction on the function f , f can be *tailored* to the sample E and thus perforce eliminate \mathbf{H} even if \mathbf{H} obtains. What needs to be precluded, then, is the tailoring of f to E . Alternatively, f needs to be independent (in some appropriate sense) of the sample E .

Let us now return to the question left hanging earlier: In eliminating \mathbf{H} , what is so special about basing the extremal sets T^γ and T_δ on the probability density function f associated with the chance hypothesis \mathbf{H} (that is, \mathbf{H} induces the probability measure $\mathbf{P}(\cdot|\mathbf{H})$ that can be represented as $f d\mathbf{U}$)? Answer: There is nothing special about f being the probability density function associated with \mathbf{H} ; instead, what is important is that f be capable of being defined independently of E , the event or sample that is observed. And indeed, Fisher's approach to eliminating chance hypotheses has already been extended in this way, though the extension, thus far, has mainly been tacit rather than explicit.

4 Specifications via Compressibility

As an example of extremal sets that, though not induced by probability density functions, nonetheless convincingly eliminate chance within Fisher's approach to hypothesis testing, let us turn to the work of Gregory Chaitin, Andrei Kolmogorov, and Ray Solomonoff. In the 1960s, Chaitin, Kolmogorov, and Solomonoff investigated what makes a sequence of coin flips random.¹⁶ Also known as *algorithmic information theory*, the Chaitin-Kolmogorov-Solomonoff theory began by noting that conventional probability theory is incapable of distinguishing bit strings of identical length (if we think of bit strings as sequences of coin tosses, then any sequence of n flips has probability 1 in 2^n).

Consider a concrete case. If we flip a fair coin and note the occurrences of heads and tails in order, denoting heads by 1 and tails by 0, then a sequence of 100 coin flips looks as follows:

defined as the length of the shortest program u (i.e., input sequence) such that $\varphi(u) = v$ (recall that the length of u is N minus the number of uninterrupted 0s at the end of u); if no such u exists (i.e., if there is no u that φ maps onto v), then we define $f(v) = N$. The function f is integer-valued and ranges between 0 and N . Moreover, given a real number δ , it induces the following extremal set (the reference class Ω here comprises all possible sequences of N bits in the output memory-register):

$$T_\delta = \{v \in \Omega \mid f(v) \leq \delta\}.$$

As a matter of simple combinatorics, it now follows that if we take δ to be an integer between 0 and N , the cardinality of T_δ (i.e., the number of elements in T_δ) is no greater than $2^{\delta+1}$. If we denote the cardinality of T_δ by $|T_\delta|$, this means that $|T_\delta| \leq 2^{\delta+1}$. The argument demonstrating this claim is straightforward: The function φ that maps input sequences to output sequences associates at most one output sequence to any given input sequence (which is not to say that the same output sequence may not be mapped onto by many input sequences). Since for any integer δ between 0 and N , there are at most 1 input sequence of length 0, 2 input sequences of length 1, 4 input sequences of length 2, ..., and 2^δ input sequences of length δ , it follows that $|T_\delta| \leq 1 + 2 + 4 + \dots + 2^\delta = 2^{\delta+1} - 1 < 2^{\delta+1}$.

Suppose now that \mathbf{H} is a chance hypothesis characterizing the tossing of a fair coin. Any output sequence v in the reference class Ω will therefore have probability 2^{-N} . Moreover, since the extremal set T_δ contains at most $2^{\delta+1}$ elements of Ω , it follows that the probability of the extremal set T_δ conditional on \mathbf{H} will be bounded as follows:

$$\mathbf{P}(T_\delta | \mathbf{H}) \leq 2^{\delta+1} / 2^N = 2^{\delta+1-N}.$$

For N large and δ small, this probability will be minuscule, and certainly smaller than any significance level we might happen to set. Consequently, for the sequences with short programs (i.e., those whose programs have length no greater than δ), Fisher's approach applied to such rejection regions would warrant eliminating the chance hypothesis \mathbf{H} . And indeed, this is exactly the conclusion reached by Chaitin, Kolmogorov, and Solomonoff. Kolmogorov even invoked the language of statistical mechanics to describe this result, calling the random sequences high entropy sequences, and the nonrandom sequence low entropy sequences.¹⁹ To sum up, the collection of algorithmically compressible (and therefore nonrandom) sequences has small probability among the totality of sequences, so that observing such a sequence is reason to look for explanations other than chance.

5 Prespecifications vs. Specifications

There is an important distinction between prespecifications and specifications that I now need to address. To do so, let me review where we are in the argument. The basic intuition I am trying to formalize is that specifications are patterns delineating events of small probability whose

occurrence cannot reasonably be attributed to chance. With such patterns, we saw that the order in which pattern and event are identified can be important. True, we did see some clear instances of patterns being identified after the occurrence of events and yet being convincingly used to preclude chance in the explanation of those events (cf. the rejection regions induced by probability density functions as well as classes of highly compressible bit strings — see sections 3 and 4 respectively). Even so, for such after-the-event patterns, some additional restrictions needed to be placed on the patterns to ensure that they would convincingly eliminate chance.

In contrast, we saw that before-the-event patterns, which we called prespecifications, require no such restrictions. Indeed, by identifying a pattern in advance of an event, a prespecification precludes the identification of the event's underlying pattern with reference to its occurrence. In other words, with prespecifications you can't just read the pattern off the event. The case of prespecifications is therefore particularly simple. Not only is the pattern identified in advance of the event, but with prespecifications only one pattern corresponding to one possible event is under consideration.

To see that prespecifications place no restriction on the patterns capable of ruling out chance, consider the following bit string:

(R) 1100001101011000110111111010001100011011001110111
00011001000010111101110110011111010010100101011110.

This string in fact resulted from tossing a fair coin and treating heads as “1” and tails as “0” (we saw this string earlier in section 4). Now, imagine the following scenario. Suppose you just tossed a coin 100 times and observed this sequence. In addition, suppose a friend was watching, asked for the coin, and you gave it to him. Finally, suppose the next day you meet your friend and he remarks, “It’s the craziest thing, but last night I was tossing that coin you gave me and I observed the exact same sequence of coin tosses you observed yesterday [i.e., the sequence (R)].” After some questioning on your part, he assures you that he was in fact tossing that very coin you gave him, giving it good jolts, trying to keep things random, and in no way cheating. Question: Do you believe your friend? The natural reaction is to be suspicious and think that somewhere, somehow, your friend (or perhaps some other trickster) was playing shenanigans. Why? Because prespecified events of small probability are very difficult to recreate by chance. It’s one thing for highly improbable chance events to happen once. But for them to happen twice is just too unlikely. The intuition here is the widely accepted folk-wisdom that says “lightning doesn’t strike twice in the same place.” Note that this intuition is taken quite seriously in the sciences. It is, for instance, the reason origin-of-life researchers tend to see the origin of the genetic code as a one-time event. Although there are some variations, the genetic code is essentially universal. Thus, for the same genetic code to emerge twice by undirected material mechanisms would simply be too improbable.²⁰

With the sequence (R) treated as a prespecification, its chance occurrence is not in question but rather its chance *reoccurrence*. That raises the question, however, whether it is possible to rule

out the chance occurrence of an event in the absence of its prior occurrence. This question is important because specifications, as I indicated in section 1, are supposed to be patterns that nail down design and therefore that inherently lie beyond the reach of chance. Yet, the sequence (R) does not do this. Granted, when treated as a prespecification, (R) rules out the event's chance reoccurrence. But the sequence itself displays no pattern that would rule out its original occurrence by chance. Are there patterns that, if exhibited in events, would rule out their original occurrence by chance?

To see that the answer is yes, consider the following sequence (again, treating "1" as heads and "0" as tails; note that the designation ψR here is meant to suggest pseudo-randomness):

(ψR) 01000110110000010100111001011101110000000100100011
01000101011001111000100110101011110011011110111100.

As with the sequence (R), imagine that you and your friend are again together. This time, however, suppose you had no prior knowledge of (ψR) and yet your friend comes to you claiming that he tossed a coin last night and witnessed the sequence (ψR). Again, he assures you that he was tossing a fair coin, giving it good jolts, trying to keep things random, and in no way cheating. Do you believe him? This time you have no prespecification, that is, no pattern given to you in advance which, if the corresponding event occurred, would convince you that the event did not happen by chance. So, how will you determine whether (ψR) happened by chance?

One approach is to employ statistical tests for randomness. A standard trick of statistics professors with an introductory statistics class is to divide the class in two, having students in one half of the class each flip a coin 100 times, writing down the sequence of heads and tails on a slip of paper, and having students in the other half each generate purely with their minds a "random looking" string of coin tosses that mimics the tossing of a coin 100 times, also writing down the sequence of heads and tails on a slip of paper. When the students then hand in their slips of paper, it is the professor's job to sort the papers into two piles, those generated by flipping a fair coin and those concocted in the students' heads. To the amazement of the students, the statistics professor is typically able to sort the papers with 100 percent accuracy.

There is no mystery here. The statistics professor simply looks for a repetition of six or seven heads or tails in a row to distinguish the truly random from the pseudo-random sequences. In a hundred coin flips, one is quite likely to see six or seven such repetitions.²¹ On the other hand, people concocting pseudo-random sequences with their minds tend to alternate between heads and tails too frequently. Whereas with a truly random sequence of coin tosses there is a 50 percent chance that one toss will differ from the next, as a matter of human psychology people expect that one toss will differ from the next around 70 percent of the time. If you will, after three or four repetitions, humans trying to mimic coin tossing with their minds tend to think it's time for a change whereas coins being tossed at random suffer no such misconception.

How, then, will our statistics professor fare when confronted with the sequences (R) and (ψ R)? We know that (R) resulted from chance because it represents an actual sequence of coin tosses. What about (ψ R)? Will it be attributed to chance or to the musings of someone trying to mimic chance? According to the professor's crude randomness checker, both (R) and (ψ R) would be assigned to the pile of sequences presumed to be truly random because both contain a seven-fold repetition (seven heads in a row for (R), seven tails in a row for (ψ R)). Everything that at first blush would lead us to regard these sequences as truly random checks out. There are approximately 50 alternations between heads and tails (as opposed to the 70 that would be expected from humans trying to mimic chance). What's more, the relative frequencies of heads and tails check out (approximately 50 heads and 50 tails). Thus, it is not as though the coin supposedly responsible for generating (ψ R) was heavily biased in favor of one side versus the other.

And yet (ψ R) is anything but random. To see this, rewrite this sequence by inserting vertical strokes as follows:

(ψ R) 0|1|00|01|10|11|000|001|010|011|100|101|110|111|0000|0001|0010|0011|
 0100|0101|0110|0111|1000|1001|1010|1011|1100|1101|1110|1111|00.

By dividing (ψ R) this way it becomes evident that this sequence was constructed simply by writing binary numbers in ascending lexicographic order, starting with the one-digit binary numbers (i.e., 0 and 1), proceeding to the two-digit binary numbers (i.e., 00, 01, 10, and 11), and continuing until 100 digits were recorded. The sequence (ψ R), when continued indefinitely, is known as the Champernowne sequence and has the property that any N -digit combination of bits appears in this sequence with limiting frequency 2^{-N} . D. G. Champernowne identified this sequence back in 1933.²²

The key to defining specifications and distinguishing them from prespecifications lies in understanding the difference between sequences such as (R) and (ψ R). The coin tossing events signified by (R) and (ψ R) are each highly improbable. And yet (R), for all we know, could have, and did, arise by chance whereas (ψ R) cannot plausibly be attributed to chance. Let us now analyze why that is the case.

6 Specificity

The crucial difference between (R) and (ψ R) is that (ψ R) exhibits a simple, easily described pattern whereas (R) does not. To describe (ψ R), it is enough to note that this sequence lists binary numbers in increasing order. By contrast, (R) cannot, so far as we can tell, be described any more simply than by repeating the sequence. Thus, what makes the pattern exhibited by (ψ R) a specification is that the pattern is easily described but the event it denotes is highly improbable and therefore very difficult to reproduce by chance. It's this combination of pattern-simplicity (i.e., easy description of pattern) and event-complexity (i.e., difficulty of reproducing

the corresponding event by chance) that makes the pattern exhibited by (ψR) — but not (R) — a specification.²³

This intuitive characterization of specification needs now to be formalized. We begin with an agent S trying to determine whether an event E that has occurred did so by chance according to some chance hypothesis \mathbf{H} (or, equivalently, according to some probability distribution $\mathbf{P}(\cdot|\mathbf{H})$). S notices that E exhibits a pattern T . For simplicity, we can think of T as itself an event that is entailed by the event E . Thus, the event E might be a die toss that lands six and T might be the composite event consisting of all die tosses that land on an even face. When T is conceived as an event, we'll refer to it as a *target*. Alternatively, T can be conceived abstractly as a pattern that precisely identifies the event (target) T . In that case, we refer to T as a *pattern*. As a pattern, T is typically describable within some communication system. In that case, we may also refer to T , described in this communication system, as a *description*. Thus, for instance, the phrase “displaying an even face” describes a pattern that maps onto a (composite) event in which a die lands either two, four, or six. Typically, when describing the pattern T , I'll put the description in quotes; yet, when talking about the event that this pattern maps onto, I'll refer to it directly and leave off the quotes.

To continue our story, given that S has noticed that E exhibits the pattern T , S 's background knowledge now induces a *descriptive complexity* of T , which measures the simplest way S has of describing T . Let us denote this descriptive complexity of T by $\phi'_S(T)$.²⁴ Note that the pattern T need not be uniquely associated with the event E — S is trying to find a simply described pattern to which E conforms, but it may be that there are still simpler ones out there. Note also that simplicity here is with respect to S 's background knowledge and is therefore denoted by a complexity measure ϕ'_S , thereby making the dependence of this measure on S explicit.²⁵

How should we think of ϕ'_S ? S , to identify a pattern T exhibited by an event E , formulates a description of that pattern. To formulate such a description, S employs a communication system, that is, a system of signs. S is therefore not merely an agent but a *semiotic agent*. Accordingly, $\phi'_S(T)$ denotes the complexity/semiotic cost that S must overcome in formulating a description of the pattern T . This complexity measure is well-defined and is not rendered multivalent on account of any semiotic subsystems that S may employ. We may imagine, for instance, a semiotic agent S who has facility with various languages (let us say both natural and artificial [e.g., computational programming] languages) and is able to describe the pattern T more simply in one such language than another. In that case, the complexity $\phi'_S(T)$ will reflect the simplest of S 's semiotic options for describing T .

Since the semiotic agents that formulate patterns to eliminate chance are finite rational agents embodied in a finite dimensional spacetime manifold (i.e., our universe), there are at most countably many of them, and these in turn can formulate at most finitely many patterns. Accordingly, the totality of patterns that S 's cognitive apparatus is able to distinguish is finite; and for all such rational agents embodied in the finite dimensional spacetime manifold that is our universe (whether it is bounded or unbounded), we can represent the totality of patterns available

terms of compressibility (see section 4), most bit sequences are maximally complex. Accordingly, φ_s , when evaluated with respect to the pattern in (R), will, in all likelihood, yield a complexity on the order of 10^{30} .

For a less artificial example of specificational resources in action, imagine a dictionary of 100,000 ($= 10^5$) basic concepts. There are then 10^5 1-level concepts, 10^{10} 2-level concepts, 10^{15} 3-level concepts, and so on. If “bidirectional,” “rotary,” “motor-driven,” and “propeller” are basic concepts, then the molecular machine known as the bacterial flagellum can be characterized as a 4-level concept of the form “bidirectional rotary motor-driven propeller.” Now, there are approximately $N = 10^{20}$ concepts of level 4 or less, which therefore constitute the specificational resources relevant to characterizing the bacterial flagellum. Next, define $p = \mathbf{P}(T|\mathbf{H})$ as the probability for the chance formation for the bacterial flagellum. T , here, is conceived not as a pattern but as the evolutionary event/pathway that brings about that pattern (i.e., the bacterial flagellar structure). Moreover, \mathbf{H} , here, is the relevant chance hypothesis that takes into account Darwinian and other material mechanisms. We may therefore think of the specificational resources as allowing as many as $N = 10^{20}$ possible targets for the chance formation of the bacterial flagellum, where the probability of hitting each target is not more than p . Factoring in these N specificational resources then amounts to checking whether the probability of hitting any of these targets by chance is small, which in turn amounts to showing that the product Np is small.

The negative logarithm to the base 2 of this last number, $-\log_2 Np$, we now define as the *specificity* of the pattern in question. Thus, for a pattern T , a chance hypothesis \mathbf{H} , and a semiotic agent S for whom φ_s measures specificational resources, the specificity σ is given as follows:

$$\sigma = -\log_2[\varphi_s(T) \cdot \mathbf{P}(T|\mathbf{H})].$$

Note that T in $\varphi_s(T)$ is treated as a pattern and that T in $\mathbf{P}(T|\mathbf{H})$ is treated as an event (i.e., the event identified by the pattern).

What is the meaning of this number, the specificity σ ? To unpack σ , consider first that the product $\varphi_s(T) \cdot \mathbf{P}(T|\mathbf{H})$ provides an upper bound on the probability (with respect to the chance hypothesis \mathbf{H}) for the chance occurrence of an event that matches any pattern whose descriptive complexity is no more than T and whose probability is no more than $\mathbf{P}(T|\mathbf{H})$. The intuition here is this: think of S as trying to determine whether an archer, who has just shot an arrow at a large wall, happened to hit a tiny target on that wall by chance. The arrow, let us say, is indeed sticking squarely in this tiny target. The problem, however, is that there are lots of other tiny targets on the wall. Once all those other targets are factored in, is it still unlikely that the archer could have hit any of them by chance? That’s what $\varphi_s(T) \cdot \mathbf{P}(T|\mathbf{H})$ computes, namely, whether of all the other targets \tilde{T} for which $\mathbf{P}(\tilde{T}|\mathbf{H}) \leq \mathbf{P}(T|\mathbf{H})$ and $\varphi_s(\tilde{T}) \leq \varphi_s(T)$, the probability of any of these targets being hit by chance according to \mathbf{H} is still small. These other targets \tilde{T} are ones that, in other circumstances, S might have picked to match up an observed event with a pattern of equal or lower descriptive complexity than T . The additional requirement that these other targets

\tilde{T} have probability no more than $\mathbf{P}(T|\mathbf{H})$ ensures that S is ruling out large targets in assessing whether E happened by chance. Hitting large targets by chance is not a problem. Hitting small targets by chance can be.

We may therefore think of $\varphi_S(T) \cdot \mathbf{P}(T|\mathbf{H})$ as gauging the degree to which S might have been self-consciously adapting the pattern T to the observed event E rather than allowing the pattern simply to flow out of the event. Alternatively, we may think of this product as providing a measure of the artificiality of imposing the pattern T on E . For descriptively simple patterns whose corresponding target has small probability, the artificiality is minimized. Note that putting the logarithm to the base 2 in front of the product $\varphi_S(T) \cdot \mathbf{P}(T|\mathbf{H})$ has the effect of changing scale and directionality, turning probabilities into number of bits and thereby making the specificity a measure of information.²⁷ This logarithmic transformation therefore ensures that the simpler the patterns and the smaller the probability of the targets they constrain, the larger specificity.

To see that the specificity so defined corresponds to our intuitions about specificity in general, think of the game of poker and consider the following three descriptions of poker hands: “single pair,” “full house,” and “royal flush.” If we think of these poker hands as patterns denoted respectively by T_1 , T_2 , and T_3 , then, given that they each have the same description length (i.e., two words for each), it makes sense to think of these patterns as associated with roughly equal specificational resources. Thus, $\varphi_S(T_1)$, $\varphi_S(T_2)$, and $\varphi_S(T_3)$ will each be roughly equal. And yet, the probability of one pair far exceeds the probability of a full house which, in turn, far exceeds the probability of a royal flush. Indeed, there are only 4 distinct royal-flush hands but 3744 distinct full-house hands and 1,098,240 distinct single-pair hands,²⁸ implying that $\mathbf{P}(T_2|\mathbf{H})$ will be three orders of magnitude larger than $\mathbf{P}(T_3|\mathbf{H})$ and $\mathbf{P}(T_1|\mathbf{H})$ will be three orders of magnitude larger than $\mathbf{P}(T_2|\mathbf{H})$ (T_1 , T_2 , and T_3 corresponding respectively to a single pair, a full house, and a royal flush, and \mathbf{H} corresponding to random shuffling). It follows that $\varphi_S(T_2) \cdot \mathbf{P}(T_2|\mathbf{H})$ will be three orders of magnitude larger than $\varphi_S(T_3) \cdot \mathbf{P}(T_3|\mathbf{H})$ (since we are assuming that $\varphi_S(T_2)$ and $\varphi_S(T_3)$ are roughly equal), implying that when we take the negative logarithm to the base 2, the specificity associated with the full house pattern will be about 10 less than the specificity associated with the royal flush pattern. Likewise, the specificity of the single pair pattern will be about 10 less than that of the full house pattern.

In this example, specificational resources were roughly equal and so they could be factored out. But consider the following description of a poker hand: “four aces and the king of diamonds.” For the underlying pattern here, which we’ll denote by T_4 , this description is about as simple as it can be made. Since there is precisely one poker hand that conforms to this description, its probability will be one-fourth the probability of getting a royal flush, i.e., $\mathbf{P}(T_4|\mathbf{H}) = \mathbf{P}(T_3|\mathbf{H})/4$. And yet, $\varphi_S(T_4)$ will be a lot bigger than $\varphi_S(T_3)$. Note that φ_S is not counting the number of words in the description of a pattern but the number of patterns of comparable or lesser complexity in terms of word count, which will make $\varphi_S(T_4)$ many orders of magnitude bigger than $\varphi_S(T_3)$. Accordingly, even though $\mathbf{P}(T_4|\mathbf{H})$ and $\mathbf{P}(T_3|\mathbf{H})$ are pretty close, $\varphi_S(T_4) \cdot \mathbf{P}(T_4|\mathbf{H})$ will be many orders of magnitude bigger than $\varphi_S(T_3) \cdot \mathbf{P}(T_3|\mathbf{H})$, implying that the specificity associated with T_4 will be substantially smaller than the specificity associated with T_3 . This may seem

counterintuitive because, in absolute terms, “four aces and the king of diamonds” more precisely circumscribes a poker hand than “royal flush” (with the former, there is precisely one hand, whereas with the latter, there are four). Indeed, we can define the absolute specificity of T as $-\log_2 \mathbf{P}(T|\mathbf{H})$. But specificity, as I’m defining it here, includes not just absolute specificity but also the cost of describing the pattern in question. Once this cost is included, the specificity of “royal flush” exceeds that of “four aces and the king of diamonds.”

7 Specified Complexity

Let us now return to our point of departure, namely, an agent S trying to show that an event E that has occurred is not properly attributed to a chance hypothesis \mathbf{H} . Suppose that E conforms to the pattern T and that T has high specificity, that is, $-\log_2[\varphi_S(T) \cdot \mathbf{P}(T|\mathbf{H})]$ is large or, correspondingly, $\varphi_S(T) \cdot \mathbf{P}(T|\mathbf{H})$ is positive and close to zero. Is this enough to show that E did not happen by chance? No. What specificity tells us is that a single archer with a single arrow is less likely than not (i.e., with probability less than 1/2) to hit the totality of targets whose probability is less than or equal to $\mathbf{P}(T|\mathbf{H})$ and whose corresponding patterns have descriptive complexity less than or equal to that of T . But what if there are multiple archers shooting multiple arrows? Given enough archers with enough arrows, even if the probability of any one of them hitting a target with a single arrow is bounded above by $\varphi_S(T) \cdot \mathbf{P}(T|\mathbf{H})$, and is therefore small, the probability of at least one of them hitting such a target with numerous arrows may be quite large. It depends on how many archers and how many arrows are available.

More formally, if a pattern T is going to be adequate for eliminating the chance occurrence of E , it is not enough just to factor in the probability of T and the specificational resources associated with T . In addition, we need to factor in what I call the *replicational resources* associated with T , that is, all the opportunities to bring about an event of T ’s descriptive complexity and improbability by multiple agents witnessing multiple events. If you will, the specificity $\varphi_S(T) \cdot \mathbf{P}(T|\mathbf{H})$ (sans negative logarithm) needs to be supplemented by factors M and N where M is the number of semiotic agents (cf. archers) that within a context of inquiry might also be witnessing events and N is the number of opportunities for such events to happen (cf. arrows). Just because a single archer shooting a single arrow may be unlikely to hit one of several tiny targets, once the number of archers M and the number of arrows N are factored in, it may nonetheless be quite likely that some archer shooting some arrow will hit one of those targets. As it turns out, the probability of some archer shooting some arrow hitting some target is bounded above by

$$M \cdot N \cdot \varphi_S(T) \cdot \mathbf{P}(T|\mathbf{H})$$

If, therefore, this number is small (certainly less than 1/2 and preferably close to zero), it follows that it is less likely than not for an event E that conforms to the pattern T to have happened according to the chance hypothesis \mathbf{H} . Simply put, if \mathbf{H} was operative in the production of some event in S ’s context of inquiry, something other than E should have happened even with all the

replicational and specificational resources relevant to E 's occurrence being factored in. Note that we refer to replicational and specificational resources jointly as *probabilistic resources* (compare section 2). Moreover, we define the logarithm to the base 2 of $M \cdot N \cdot \varphi_S(T) \cdot \mathbf{P}(T|\mathbf{H})$ as the *context-dependent specified complexity* of T given \mathbf{H} , the context being S 's context of inquiry:

$$\tilde{\chi} = -\log_2[M \cdot N \cdot \varphi_S(T) \cdot \mathbf{P}(T|\mathbf{H})].$$

Note that the tilde above the Greek letter chi indicates $\tilde{\chi}$'s dependence on the replicational resources within S 's context of inquiry. In a moment, we'll consider a form of specified complexity that is independent of the replicational resources associated with S 's context of inquiry and thus, in effect, independent of S 's context of inquiry period (thereby strengthening the elimination of chance and the inference to design). For most purposes, however, $\tilde{\chi}$ is adequate for assessing whether T happened by chance. The crucial cut-off, here, is $M \cdot N \cdot \varphi_S(T) \cdot \mathbf{P}(T|\mathbf{H}) < 1/2$: in this case, the probability of T happening according to \mathbf{H} given that all relevant probabilistic resources are factored is strictly less than 1/2, which is equivalent to $\tilde{\chi} = -\log_2[M \cdot N \cdot \varphi_S(T) \cdot \mathbf{P}(T|\mathbf{H})]$ being strictly greater than 1. Thus, if $\tilde{\chi} > 1$, it is less likely than not that an event of T 's descriptive complexity and improbability would happen according to \mathbf{H} even if as many probabilistic resources as are relevant to T 's occurrence are factored in.

To see how all this works, consider the following example from Dan Brown's wildly popular novel *The Da Vinci Code*. The heroes, Robert Langdon and Sophie Neveu, find themselves in an ultra-secure, completely automated portion of a Swiss bank ("the Depository Bank of Zurich"). Sophie's grandfather, before dying, had revealed the following ten digits separated by hyphens: 13-3-2-21-1-1-8-5. Langdon is convinced that this is the bank account number that will open a deposit box containing crucial information about the Holy Grail. We pick up the storyline here:²⁹

The cursor blinked. Waiting.

Ten digits. Sophie read the numbers off the printout, and Langdon typed them in.

ACCOUNT NUMBER:
1332211185

When he had typed the last digit, the screen refreshed again. A message in several languages appeared. English was on top.

CAUTION:
Before you strike the enter key, please check the
accuracy of your account number.
For your own security, if the computer does not
recognize the account number, this system will
automatically shut down.

"*Fonction terminer*," Sophie said, frowning. "Looks like we only get one try." Standard ATM machines allowed users three attempts to type in a PIN before confiscating their bank card. This was obviously no ordinary cash machine....

"No." She pulled her hand away. "This isn't the right account number."

“Of course it is! Ten digits. What else would it be?”

“It’s too random.”

Too random? Langdon could not have disagreed more. Every bank advised its customers to choose PINs at random so nobody could guess them. Certainly clients *here* would be advised to choose their account numbers at random.

Sophie deleted everything they had just typed in and looked up at Langdon, her gaze self-assured. “It’s far too coincidental that this supposedly random account number could be rearranged to form the Fibonacci sequence.”

[The digits that Sophie’s grandfather made sure she received posthumously, namely, 13-3-2-21-1-1-8-5, can be rearranged as 1-1-2-3-5-8-13-21, which are the first eight numbers in the famous Fibonacci sequence. In this sequence, numbers are formed by adding the two immediately preceding numbers. The Fibonacci sequence has some interesting mathematical properties and even has applications to biology.³⁰]

Langdon realized she had a point. Earlier, Sophie had rearranged this account number into the Fibonacci sequence. What were the odds of being able to do that?

Sophie was at the keypad again, entering a different number, as if from memory. “Moreover, with my grandfather’s love of symbolism and codes, it seems to follow that he would have chosen an account number that had meaning to him, something he could easily remember.” She finished typing the entry and gave a sly smile. “Something that appeared random but was not.”

Needless to say, Robert and Sophie punch in the Fibonacci sequence 1123581321 and retrieve the crucial information they are seeking.

This example offers several lessons about the (context-dependent) specified complexity $\tilde{\chi}$. Let’s begin with $\mathbf{P}(T|\mathbf{H})$, T here is the Fibonacci sequence 1123581321. This sequence, if produced at random (i.e., with respect to the uniform probability distribution denoted by \mathbf{H}), would have probability 10^{-10} , or 1 in 10 billion. This is $\mathbf{P}(T|\mathbf{H})$. Note that for typical ATM cards, there are usually sixteen digits, and so the probability is typically on the order of 10^{-15} (not 10^{-16} because the first digit is usually fixed; for instance, Visa cards all begin with the digit “4”).

What about the specificational resources associated with 1123581321, that is, $\varphi_S(T)$? S here is either Sophie Neveu or Robert Langdon. Because T is the Fibonacci sequence, $\varphi_S(T)$, when computed according to an appropriate minimal description length metric, will be considerably less than what it would be for a truly random sequence (for a truly random sequence, $\varphi_S(T)$ would be on the order of 10^{10} for most human inquirers S — recall section 4). What about M and N ? In this context, M is the number of people who might enter the Depository Bank of Zurich to retrieve the contents of an account. Given the exclusiveness of this bank as represented in the novel, it would be fair to say that not more than 100 people (allowing repeated visits by a given person) would attempt to do this per day. Over the course of 100 years, or 36,500 days, there would in consequence not be more than 3,650,000 attempts to retrieve information from any accounts at the Depository Bank of Zurich over its entire history. That leaves N . N is the number of attempts each such person has to retrieve the contents of an account. For most ATM machines, N is 3. But for the ultra-secure Depository Bank of Zurich, N is 1. Indeed, the smaller N , the better the security. Thus, with the *Da Vinci Code* example,

$$\tilde{\chi} = -\log_2[3,650,000 \cdot \varphi_S(T) \cdot 10^{-10}] \approx -\log_2[10^{-4} \cdot \varphi_S(T)].$$

In general, the bigger $M \cdot N \cdot \varphi_S(T) \cdot \mathbf{P}(T|\mathbf{H})$ — and, correspondingly, the smaller its negative logarithm (i.e., $\tilde{\chi}$) — the more plausible it is that the event denoted by T could happen by chance. In this case, if T were a truly random sequence (as opposed to the Fibonacci sequence), $\varphi_S(T)$, as we observed earlier, would be on the order of 10^{10} , so that

$$\tilde{\chi} \approx -\log_2[10^{-4} \cdot \varphi_S(T)] \approx -\log_2[10^{-4} \cdot 10^{10}] \approx -20.$$

On the other hand, if $\varphi_S(T)$ were on the order of 10^3 or less, $\tilde{\chi}$ would be greater than 1, which would suggest that chance should be eliminated. The chance event, in this case, would be that the bank account number had been accidentally reproduced and its contents accidentally retrieved, and so the elimination of chance would amount to the inference that this had not happened accidentally. Clearly, reproducing the right account number did not happen by accident in the case of Neveu and Langdon — they were given specific information about the account from Neveu’s grandfather. Even so, this example points up the importance, in general, of password security and choosing passwords that are random. A value of $\tilde{\chi}$ bigger than 1 would point to the probability $\mathbf{P}(T|\mathbf{H})$ swamping probabilistic resources $M \cdot N \cdot \varphi_S(T)$ (i.e., for the order of magnitude of $1/\mathbf{P}(T|\mathbf{H})$ to be considerably greater than that of $M \cdot N \cdot \varphi_S(T)$), so that if T happened, it could not properly be referred to chance.

Another way to think about this example is from the Depository Bank of Zurich’s perspective. This bank wants only customers with specific information about their accounts to be able to access those accounts. It does not want accounts to be accessed by chance. Hence, it does not want the probabilistic resources $M \cdot N \cdot \varphi_S(T)$ to swamp the probability $\mathbf{P}(T|\mathbf{H})$ (i.e., for the order of magnitude of $M \cdot N \cdot \varphi_S(T)$ to exceed $1/\mathbf{P}(T|\mathbf{H})$). If that were to be the case, accounts would be in constant danger of being randomly accessed. This is why the Depository Bank of Zurich required a 10-digit account number and why credit cards employ 16-digit account numbers (along with expiration dates and 3-digit security codes on the back). If, for instance, account numbers were limited to three digits, there would be at most 1,000 different account numbers, and so, with millions of users, it would be routine that accounts would be accessed accidentally (i.e., $\mathbf{P}(T|\mathbf{H})$ would equal a mere $1/1,000$ and get swamped by $M \cdot N \cdot \varphi_S(T)$, which would be in the millions). Bottom line: the (context-dependent) specified complexity $\tilde{\chi}$ is the key in all such assessments of account security. Indeed, it provides a precise determination and measurement of account security.

As defined, $\tilde{\chi}$ is context sensitive, tied to the background knowledge of a semiotic agent S and to the context of inquiry within which S operates. Even so, it is possible to define specified complexity so that it is not context sensitive in this way. Theoretical computer scientist Seth Lloyd has shown that 10^{120} constitutes the maximal number of bit operations that the known, observable universe could have performed throughout its entire multi-billion year history.³¹ This number sets an upper limit on the number of agents that can be embodied in the universe and the

number of events that, in principle, they can observe. Accordingly, for any context of inquiry in which S might be endeavoring to determine whether an event that conforms to a pattern T happened by chance, $M \cdot N$ will be bounded above by 10^{120} . We thus define the *specified complexity* of T given \mathbf{H} (minus the tilde and context sensitivity) as

$$\chi = -\log_2[10^{120} \cdot \varphi_S(T) \cdot \mathbf{P}(T|\mathbf{H})].$$

To see that χ is independent of S 's context of inquiry, it is enough to note two things: (1) there is never any need to consider replicational resources $M \cdot N$ that exceed 10^{120} (say, by invoking inflationary cosmologies or quantum many-worlds) because to do so leads to a wholesale breakdown in statistical reasoning, and that's something no one in his saner moments is prepared to do (for the details about the fallacy of inflating one's replicational resources beyond the limits of the known, observable universe, see my article "The Chance of the Gaps"³²). (2) Even though χ depends on S 's background knowledge through $\varphi_S(T)$, and therefore appears still to retain a subjective element, the elimination of chance only requires a single semiotic agent who has discovered the pattern in an event that unmask its non-chance nature. Recall the Champernowne sequence discussed in sections 5 and 6 (i.e., $(\psi\mathbf{R})$). It doesn't matter if you are the only semiotic agent in the entire universe who has discovered its binary-numerical structure. That discovery is itself an objective fact about the world, and it rightly gets incorporated into χ via $\varphi_S(T)$. Accordingly, that sequence would not rightly be attributed to chance precisely because you were the one person in the universe to appreciate its structure.

It follows that if $10^{120} \cdot \varphi_S(T) \cdot \mathbf{P}(T|\mathbf{H}) < 1/2$ or, equivalently, that if $\chi = -\log_2[10^{120} \cdot \varphi_S(T) \cdot \mathbf{P}(T|\mathbf{H})] > 1$, then it is less likely than not on the scale of the whole universe, with all replicational and specificational resources factored in, that E should have occurred according to the chance hypothesis \mathbf{H} . Consequently, we should think that E occurred by some process other than one characterized by \mathbf{H} . Since specifications are those patterns that are supposed to underwrite a design inference, they need, minimally, to entitle us to eliminate chance. Since to do so, it must be the case that

$$\chi = -\log_2[10^{120} \cdot \varphi_S(T) \cdot \mathbf{P}(T|\mathbf{H})] > 1,$$

we therefore define *specifications* as any patterns T that satisfy this inequality. In other words, specifications are those patterns whose specified complexity is strictly greater than 1. Note that this definition automatically implies a parallel definition for *context-dependent specifications*: a pattern T is a context-dependent specification provided that its context-dependent specified complexity $\tilde{\chi}$ is strictly greater than 1. Such context-dependent specifications are widely employed in adjudicating between chance and design (cf. the Da Vinci Code example). Yet, to be secure in eliminating chance and inferring design on the scale of the universe, we need the context-independent form of specification.

As an example of specification and specified complexity in their context-independent form, let us return to the bacterial flagellum. Recall the following description of the bacterial flagellum

given in section 6: “bidirectional rotary motor-driven propeller.” This description corresponds to a pattern T . Moreover, given a natural language (English) lexicon with 100,000 ($= 10^5$) basic concepts (which is supremely generous given that no English speaker is known to have so extensive a basic vocabulary), we estimated the complexity of this pattern at approximately $\phi_S(T) = 10^{20}$ (for definiteness, let’s say S here is me; any native English speaker with a some of knowledge of biology and the flagellum would do). It follows that $-\log_2[10^{120} \cdot \phi_S(T) \cdot \mathbf{P}(T|\mathbf{H})] > 1$ if and only if $\mathbf{P}(T|\mathbf{H}) < \frac{1}{2} \times 10^{-140}$, where \mathbf{H} , as we noted in section 6, is an evolutionary chance hypothesis that takes into account Darwinian and other material mechanisms and T , conceived not as a pattern but as an event, is the evolutionary pathway that brings about the flagellar structure (for definiteness, let’s say the flagellar structure in *E. coli*). Is $\mathbf{P}(T|\mathbf{H})$ in fact less than $\frac{1}{2} \times 10^{-140}$, thus making T a specification? The precise calculation of $\mathbf{P}(T|\mathbf{H})$ has yet to be done. But some methods for decomposing this probability into a product of more manageable probabilities as well as some initial estimates for these probabilities are now in place.³³ These preliminary indicators point to T ’s specified complexity being greater than 1 and to T in fact constituting a specification.

8 Design Detection

Having defined specification, I want next to show how this concept works in eliminating chance and inferring design. Inferring design by eliminating chance is an old problem. Almost 300-years ago, the mathematician Abraham de Moivre addressed it as follows:

The same Arguments which explode the Notion of Luck, may, on the other side, be useful in some Cases to establish a due comparison between Chance and Design: We may imagine Chance and Design to be, as it were, in Competition with each other, for the production of some sorts of Events, and may calculate what Probability there is, that those Events should be rather owing to one than to the other. To give a familiar Instance of this, Let us suppose that two Packs of Piquet-Cards being sent for, it should be perceived that there is, from Top to Bottom, the same Disposition of the Cards in both packs; let us likewise suppose that, some doubt arising about this Disposition of the Cards, it should be questioned whether it ought to be attributed to Chance, or to the Maker’s Design: In this Case the Doctrine of Combinations decides the Question; since it may be proved by its Rules, that there are the odds of above 263130830000 Millions of Millions of Millions to One, that the Cards were designedly set in the Order in which they were found.³⁴

To “explode the notion of luck,” de Moivre requires a highly improbable prespecified event in which one ordered deck of cards rules out the chance reoccurrence of another with the same order. In particular, he places chance and design in competition so that the defeat of chance on the basis of improbability leads to the victory of design. Resolving this competition between chance and design is the whole point of specification.

The fundamental claim of this paper is that for a chance hypothesis \mathbf{H} , if the specified complexity $\chi = -\log_2[10^{120} \cdot \phi_S(T) \cdot \mathbf{P}(T|\mathbf{H})]$ is greater than 1, then T is a specification and the semiotic agent S is entitled to eliminate \mathbf{H} as the explanation for the occurrence of any event E

that conforms to the pattern T (S is similarly entitled to eliminate \mathbf{H} when the context-dependent specified complexity $\tilde{\chi} = -\log_2[M \cdot N \cdot \phi_S(T) \cdot \mathbf{P}(T|\mathbf{H})]$ is greater than one, only this time, because $M \cdot N$ will be less than 10^{120} , the strength with which $\tilde{\chi}$ eliminates \mathbf{H} will be less than what it is for χ).

The first question we therefore need to answer is whether a specified complexity greater than 1 adequately eliminates the chance hypothesis \mathbf{H} at all. Bayesian and likelihood approaches say no, claiming that statistical rationality permits the elimination of chance hypotheses only in comparison with other chance hypotheses (thus treating design itself as a chance hypothesis). But these approaches, I argue in Addendum 2, invariably presuppose an account of specification. Moreover, the conceptual difficulties that proponents of comparative approaches to statistical rationality have cited against straight chance-elimination approaches are, in my view, met by the apparatus of specified complexity, a point I argued at length in *The Design Revolution*.³⁵ I will therefore take it as settled that specified complexity is an adequate tool for eliminating individual chance hypotheses.

Now, in many statistical applications, the elimination of \mathbf{H} (whether on Fisherian or Bayesian or other grounds), does not rule out chance as such but merely invites alternative chance hypotheses \mathbf{H}' . (Think of a lopsided die and \mathbf{H} as the hypothesis that all sides will appear with probability $1/6$; even if specified complexity eliminates \mathbf{H} , that still leaves alternative hypotheses \mathbf{H}' for which the probability of the faces are not all equal.) Thus, if specified complexity is to rule out chance *überhaupt*, we must have a good grasp of what chance hypotheses would have been operating to produce the observed event E whose chance status is in question. Suppose the relevant collection of chance hypotheses that we have good reason to think were operating if the event E happened by chance is some collection of chance hypotheses $\{\mathbf{H}_i\}_{i \in I}$ indexed by the index set I . Then, to eliminate all these chance hypotheses, $\chi_i = -\log_2[10^{120} \cdot \phi_S(T) \cdot \mathbf{P}(T|\mathbf{H}_i)]$ must be greater than 1 for each \mathbf{H}_i .

Granted, this would eliminate all the chance hypotheses in $\{\mathbf{H}_i\}_{i \in I}$ but would it eliminate all chance hypotheses *überhaupt*? Probabilistic arguments are inherently fallible in the sense that our assumptions about relevant probability distributions might always be in error. Thus, it is always a possibility that $\{\mathbf{H}_i\}_{i \in I}$ omits some crucial chance hypothesis that might be operating in the world and account for the event E in question. But are we to take this possibility seriously in the absence of good evidence for the operation of such a chance hypothesis in the production of E ? Indeed, the mere possibility that we might have missed some chance hypothesis is hardly reason to think that such a hypothesis was operating. Nor is it reason to be skeptical of a design inference based on specified complexity. Appealing to the unknown to undercut what we do know is never sound epistemological practice. Sure, we may be wrong. But unknown chance hypotheses (and the unknown material mechanisms that supposedly induce them) have no epistemic force in showing that we are wrong. Inquiry can throw things into question only by taking other things to be fixed. The unknown is not a fixed point. It cannot play such a role.

To see that it is illegitimate to appeal to unknown chance hypotheses to undercut the chance elimination methods that I am proposing, consider the place where these methods are most hotly contested, namely, biological origins. Take, for instance, the well known Miller-Urey experiment.³⁶ This was one of the early primitive atmospheric simulation experiments designed to show that the building blocks of life could readily have formed on the early Earth. My interest here is not in how plausibly those experiments model what happened on the early Earth, but the logic by which these experiments were interpreted and by which inferences were drawn from them. In the Miller-Urey experiment, various compounds were placed in an apparatus, zapped with sparks to simulate lightning, and then the product was collected in a trap. Lo and behold, biologically significant chemical compounds were discovered, notably certain amino acids.

In the 1950s, when this experiment was performed, it was touted as showing that a purely chemical solution to the origin of life was just around the corner. Since then, this enthusiasm has waned because such experiments merely yield certain rudimentary building blocks for life. No experiments since then have shown how these building blocks could, by purely chemical means (and thus apart from design), be built up into complex biomolecular systems needed for life (like proteins and multiprotein assemblages, to say nothing of fully functioning cells).³⁷ The Miller-Urey experiment showed that certain biologically significant chemical building blocks, which previously had only been synthesized in living systems, were now capable of being synthesized abiologically. Moreover, it showed that this could be done easily and reproducibly — anyone with the requisite laboratory equipment could replicate the experiment and its results. What this means, in statistical terms, is that the Miller-Urey experiment yielded its biologically significant chemical building blocks with high probability. Scientists, in turn, took this as evidence that a purely chemical solution of life's origin was credible and that the need to invoke design or intelligence was dispensable.

Here, again, we see de Moivre's point about chance and design being in competition, only this time chance comes out the winner because the probabilities are sufficiently large.³⁸ But if large probabilities vindicate chance and defeat design, why shouldn't small probabilities do the opposite — vindicate design and defeat chance? Indeed, in many special sciences, everything from forensics to archeology to SETI (the Search for Extraterrestrial Intelligence), small probabilities do just that. Objections only get raised against inferring design on the basis of such small probability, chance elimination arguments when the designers implicated by them are unacceptable to a materialistic worldview, as happens at the origin of life, whose designer could not be an intelligence that evolved through purely materialistic processes. Parity of reasoning demands that if large probabilities vindicate chance and defeat design, then small probabilities should vindicate design and defeat chance. The job of specified complexity is to marshal these small probabilities in a way that convincingly defeats chance and vindicates design.

At this point, critics of specified complexity raise two objections. First, they contend that because we can never know all the chance hypotheses responsible for a given outcome, to infer design because specified complexity eliminates a limited set of chance hypotheses constitutes an

argument from ignorance. But this criticism is misconceived. The argument from ignorance, also known as the *appeal to ignorance* or by the Latin *argumentum ad ignorantiam*, is

the fallacy of arguing that something must be true because nobody can prove it false or, alternatively, that something must be false because nobody can prove it true. Such arguments involve the illogical notion that one can view the lack of evidence about a proposition as being positive evidence for it or against it. But lack of evidence is lack of evidence, and supports no conclusion. An example of an appeal to ignorance: Styrofoam cups must be safe; after all, no studies have implicated them in cancer. The argument is fallacious because it is possible that no studies have been done on those cups or that what studies have been done have not focused on cancer (as opposed to other diseases).³⁹

In eliminating chance and inferring design, specified complexity is not party to an argument from ignorance. Rather, it is underwriting an *eliminative induction*. Eliminative inductions argue for the truth of a proposition by actively refuting its competitors (and not, as in arguments from ignorance, by noting that the proposition has yet to be refuted). Provided that the proposition along with its competitors form a mutually exclusive and exhaustive class, eliminating all the competitors entails that the proposition is true. (Recall Sherlock Holmes’s famous dictum: “When you have eliminated the impossible, whatever remains, however improbable, must be the truth.”) This is the ideal case, in which eliminative inductions in fact become deductions. But eliminative inductions can be convincing without knocking down every conceivable alternative, a point John Earman has argued effectively. Earman has shown that eliminative inductions are not just widely employed in the sciences but also indispensable to science.⁴⁰ Suffice it to say, by refusing the eliminative inductions by which specified complexity eliminates chance, one artificially props up chance explanations and (let the irony not be missed) *eliminates design explanations* whose designing intelligences don’t match up conveniently with a materialistic worldview.⁴¹

And this brings us to the other objection, namely, that we must know something about a designer’s nature, purposes, propensities, causal powers, and methods of implementing design before we can legitimately determine whether an object is designed. I refer to the requirement that we must have this independent knowledge of designers as the *independent knowledge requirement*. This requirement, so we are told, can be met for materially embodied intelligences but can never be met for intelligences that cannot be reduced to matter, energy, and their interactions.⁴² By contrast, to employ specified complexity to infer design is to take the view that objects, even if nothing is known about how they arose, can exhibit features that reliably signal the action of an intelligent cause. There are two competing approaches to design detection here that cut to the heart of what it is to know that something is designed. The one approach requires independent knowledge of the designer. The other says that the signature of design can be read without any such knowledge. Which approach is correct? I submit the latter, which, happily, is also consistent with employing specified complexity to infer design.

To see that the independent knowledge requirement, as a principle for deciding whether something is designed, is fundamentally misguided, consider the following admission by Elliott Sober, who otherwise happens to embrace this requirement: “To infer watchmaker from watch,

you needn't know exactly what the watchmaker had in mind; indeed, you don't even have to know that the watch is a device for measuring time. Archaeologists sometimes unearth tools of unknown function, but still reasonably draw the inference that these things are, in fact, *tools*.”⁴³ Sober's remark suggests that design inferences may look strictly to features of designed objects and thus presuppose no knowledge about the characteristics of the designer.

His remark, therefore, raises the following question: when the evidence for design is circumstantial, as invariably it is when we apply the techniques of design detection described in this paper, what do we know about any putative designer who may have been responsible for an object in question? In fact, we may know nothing about the designer. Nor would it matter. To be sure, we may know that other designers in our experience were able to bring about an object of the sort we are considering. But what if the designer actually responsible for the object brought it about by means unfathomable to us (e.g., by some undreamt of technologies)? This is the problem of *multiple realizability*, and it undercuts the independent knowledge requirement because it points up that what leads us to infer design is not knowledge of designers and their capabilities but knowledge of the patterns exhibited by designed objects (a point that specified complexity captures precisely).

This last point underscores another problem with the independent knowledge requirement, namely, what I call the problem of *inductive regress*. Suppose, like Sober, one wants to argue that independent knowledge of designers is the key to inferring design.⁴⁴ Consider now some archeologists in the field who stumble across an arrowhead. How do they know that it is indeed an arrowhead and thus the product of design? What sort of archeological background knowledge had to go into their design hypothesis? Certainly, the archeologists would need past experience with arrowheads. But how did they recognize that the arrowheads in their past experience were designed? Did they see humans actually manufacture those arrowheads? If so, how did they recognize that these humans were acting deliberately, as designing agents, and not just randomly chipping away at random chunks of rock (carpentry and sculpting entail design; but whittling and chipping, though performed by intelligent agents, do not)? As is evident from this line of reasoning, the induction needed to recognize design can never get started. Our ability to recognize design must therefore arise independently of induction and therefore independently of any independent knowledge requirement about the capacities of designers. In fact, it arises directly from the patterns in the world that signal intelligence, to wit, from specifications.

Another problem with the independent knowledge requirement is that it hinders us from inferring design that outstrips our intellectual and technological sophistication. I call this the problem of *dummied down design*: the independent knowledge requirement limits our ability to detect design to the limits we impose on designers. But such limits are artificial. Suppose, for instance, that the molecular biology of the cell is in fact intelligently designed. If so, it represents nanotechnology of a sophistication far beyond anything that human engineers are presently capable of or may ever be capable of. By the independent knowledge requirement, we have no direct experience of designers capable of such design work. Thus, even if system after molecular biological system exhibited high specified complexity, the independent evidence requirement would prevent us

from recognizing their design and keep us wedded to wholly inadequate materialistic explanations of these systems.

But consider next a thought experiment. Imagine that space travelers show up on Earth loaded with unbelievably advanced technology. They tell us (in English) that they've had this technology for hundreds of millions of years and give us solid evidence of this claim (perhaps by pointing to some star cluster hundreds of millions of light years away whose arrangement signifies a message that confirms the aliens' claim). Moreover, they demonstrate to us that with this technology they can, atom by atom and molecule by molecule, assemble the most complex organisms. Suppose we have good reason to think that these aliens were here at key moments in life's history (e.g., at the origin of life, the origin of eukaryotes, and the origin of the animal phyla in the Cambrian). Suppose, further, that in forming life from scratch the aliens would not leave any trace (their technology is so advanced that they clean up after themselves perfectly — no garbage or any other signs of activity would be left behind). Suppose, finally, that none of the facts of biology are different from what they are now. Should we now think that life at key moments in its history was designed?

We now have all the independent knowledge we could ever want for the existence and attributes of materially embodied designers capable of bringing about the complexity of life on earth. If, in addition, our best probabilistic analysis of the biological systems in question tells us that they exhibit high specified complexity and therefore that unguided material processes could not have produced them with anything like a reasonable probability, would a design inference only now be warranted? Would design, in that case, become a better explanation than materialistic evolution simply because we now have independent knowledge of designers with the capacity to produce biological systems?

This prospect, however, should raise a worry. The facts of biology, after all, have not changed, and yet design would be a better explanation if we had independent knowledge of designers capable of producing, say, the animal phyla of the Cambrian. Note that there's no smoking gun here (no direct evidence of alien involvement in the fossil record, for instance). All we know by observation is that beings with the power to generate life exist and could have acted. Would it help to know that the aliens really like building carbon-based life? But how could we know that? Do we simply take their word for it? If design is a better explanation simply because we have independent knowledge of technologically advanced space aliens, why should it not be a better explanation absent such evidence? If conventional evolutionary theory is so poor an explanation that it would cave the instant space aliens capable of generating living forms in all their complexity could be independently attested, then why should it cease to be a poor explanation absent those space aliens? The point to appreciate is that specified complexity can demonstrate this poverty of explanation even now — apart from space aliens and bizarre thought experiments.

In conclusion, suppose that our best understanding of the processes that could have been involved in the occurrence of an event *E* suggests that if *E* happened by chance, then one of the

chance hypotheses in $\{\mathbf{H}_i\}_{i \in I}$ was responsible. Suppose, further, that the specified complexity associated with each of these chance hypotheses, that is, $\chi_i = -\log_2[10^{120} \cdot \varphi_S(T) \cdot \mathbf{P}(T|\mathbf{H}_i)]$, is strictly greater than 1. In that case, we've eliminated all the chance hypotheses that might explain the occurrence of E (where E conforms to the pattern T). The inference to design is now immediate. Certainly it is a necessary condition, if a design inference is to hold, that all relevant chance hypotheses be eliminated. But this condition is also sufficient in that unknown chance hypotheses have no epistemic significance in defeating design. Moreover, the requirement that we have independent knowledge of the designer founders, as we have seen, on the multiple realizability, inductive regress, and dummied down design problems. In the presence of viable chance alternatives, design is an explanation of last resort. Yet, once specified complexity has rendered all relevant chance alternatives inviable, chance as such is eliminated and design can no longer be denied.

ACKNOWLEDGMENT: I want to thank Jay Richards for organizing a symposium on design reasoning at Calvin College back in May, 2001 as well as for encouraging me to write up my thoughts about specification that emerged from this symposium. Essentially, the last original thing I wrote about specification was in 2001 when I was preparing for this symposium and wrapping up the manuscript for my book *No Free Lunch* (which was published in 2002). Although my thinking about specification has progressed considerably in these last four years, my attentions were focused on other projects. It was therefore high time that I return to this central concept in my theory of design detection and take another pass at it. I also want to thank Rob Koons, Robin Collins, Tim and Lydia McGrew, and Del Ratzsch for their useful feedback at this symposium. Stephen Meyer and Paul Nelson were also at this symposium. They have now tracked my thoughts on specification for almost fifteen years and have been my best conversation partners on this topic. Finally, I wish to thank Richard Dawkins, whose remarks about specification and complexity in *The Blind Watchmaker* first got me thinking (back in 1989) that here was the key to eliminating chance and inferring design. As I've remarked to him in our correspondence (a correspondence that he initiated), "Thanks for all you continue to do to advance the work of intelligent design. You are an instrument in the hands of Providence however much you rail against it."

Addendum 1: Note to Readers or *TDI & NFL*

Readers familiar with my books *The Design Inference* and *No Free Lunch* will note that my treatment of specification and specified complexity there (specificity, as such, does not appear explicitly in these books, though it is there implicitly) diverges from my treatment of these concepts in this paper. The changes in my account of these concepts here should be viewed as a simplification, clarification, extension, and refinement of my previous work, not as a radical departure from it. To see this, it will help to understand what prompted this new treatment of specification and specified complexity as well as why it remains in harmony with my past treatment.

The first question a reader familiar with my past treatment of these concepts is likely to ask is whatever happened to such key notions as detachability, conditional independence, tractability, and universal probability bounds that in the past had characterized my account of specification. They are still here, but they no longer need to be made explicit because of a fundamental simplification to the account of specification that appears in this paper. The simplification results from not demanding that the theory of specification do double-duty as also a theory of prespecification. The problem, as I point out in section 5 of this paper, is that a theory of prespecification can do an excellent job of ruling out the reoccurrence of a chance events, but it cannot rule out whether prespecifications themselves resulted from chance — recall the difference between the sequence (R) and the Champernowne sequence (ψR) in sections 5 through 7. The upshot of these considerations was that a theory of prespecifications can rightly tell you that lightning won't strike twice in the same place, but it cannot tell you whether lightning that's struck only once in a given place struck there by chance.

By separating off prespecifications from specifications, the account of specifications becomes much more straightforward. With specifications, the key to overturning chance is to keep the descriptive complexity of patterns low. This was the point of the tractability condition (i.e., patterns satisfy this condition if their descriptive complexity is low). With prespecifications, by contrast, the descriptive complexity of patterns is irrelevant. Rather, with prespecifications, the key to overturning chance is the causal relation between the identification of a pattern and the occurrence of an event exhibiting the pattern. Provided that a pattern delineates an event/target of small probability and provided that an event that exhibits the pattern occurs after the pattern itself has been explicitly identified, the pattern will constitute a prespecification and therefore be suitable for eliminating chance. This was the point of the conditional independence condition (i.e., patterns automatically satisfy this condition if they are identified prior to the events that exhibit these patterns). The conditional independence condition prevents “cherry picking” or “data snooping.”

Interestingly, low descriptive complexity immediately confers conditional independence. The tractability condition, as originally conceived, ensures that the tractable patterns are those that can be constructed easily from information independent of the occurrence of the events they delineate. For prespecifications this is automatically the case because the prespecifications are

always given prior to the events in question, and thus provide the necessary information about themselves. But prespecifications need not be descriptively simple. Think of a coin that's flipped 1,000 times. The pattern it exhibits will (in all likelihood) be unique in the history of coin tossing and will not be identifiable apart from the actual event of flipping that coin 1,000 times. Such a pattern, if a prespecification, will be tractable but it will not be descriptively simple. On the other hand, a Champernowne sequence of length 1,000 can be readily constructed on the basis of a simple number-theoretic scheme. The underlying pattern here, in virtue of its descriptive simplicity, is therefore tractable with respect to information that is conditionally independent of any actual coin tossing event. In *The Design Inference*, I referred to this relation between tractability and conditional independence as *detachability* and used detachability to define specification. But given that descriptive simplicity instantly yields conditional independence, detachability itself becomes a dispensable concept once descriptive simplicity supplants tractability in the definition of specification.⁴⁵

To separate off prespecifications from specifications means that specifications can no longer be patterns of high descriptive complexity. This constitutes a slight restriction on the concept of specification (as when it used to be characterized in terms of detachability), but this restriction is well worth the clarity it brings to the concept of specification. Because specifications will no longer include prespecifications of high descriptive complexity, to rule out the chance reoccurrence of a highly improbable event whose pattern has high descriptive complexity and which was identified in advance of the event, it is necessary to modify the accounts of specificity and specified complexity for prespecifications. The key is to see that for prespecifications, the specificational resources $\varphi_s(T)$ can be taken to be simply 1. Indeed, if the pattern is specified in advance of an event, there is no point to asking what other targets of comparable descriptive complexity might happen. With prespecifications, there is one and only one target. It follows that the specificity and specified complexity (whether context-dependent or context-independent) for a target T can be defined simply by dropping $\varphi_s(T)$ from the mathematical formulas that define these concepts. Accordingly, the specificity for a prespecification T is simply $\sigma = -\log_2[\mathbf{P}(T|\mathbf{H})]$ (i.e., what we previously referred to as the “absolute specificity”) and its specified complexity is simply $\tilde{\chi} = -\log_2[M \cdot N \cdot \mathbf{P}(T|\mathbf{H})]$ (with 10^{120} substituted for $M \cdot N$ in the context-independent case).

Next, in connecting my treatment of specification here to my past work, let me address a concept that has played a central role in my previous work on specified complexity but appears to play less of a role here, namely, that of a universal probability bound. In my past work, I argued on the basis of various physical constraints on the known, observable universe that for any specification delineating a target of probability 10^{-150} , an event matching that target could not happen by chance. That argument, however, worked best for prespecifications. For specifications generally, descriptive complexity needed also to be factored in. For the examples I tended to consider, however, this never proved to be a problem because the descriptive complexity was so low that it could readily be assimilated into the universal probability bound, which was so conservative that it swamped the descriptive complexity of patterns qua specifications. My present treatment resolves this discrepancy.

Where does the universal probability bound reside in my current treatment of specification? It comes up through Seth Lloyd's 10^{120} upper limit on the number of bit operations possible in the observable universe throughout its multibillion-year history.⁴⁶ This number, since it also places an upper limit on any replicational resources $M \cdot N$, plays a key role in the definition of the context-independent formulation of specified complexity: $\chi = -\log_2[10^{120} \cdot \varphi_S(T) \cdot \mathbf{P}(T|\mathbf{H})]$. And since χ must be greater than 1 for T to be a specification, this means that $\mathbf{P}(T|\mathbf{H})$ must be less than 10^{-120} . Thus, in the treatment of specification given here, we have a universal probability bound of 10^{-120} that is thirty orders of magnitude bigger than my previous one (which sat at 10^{-150}), but for which the descriptive complexity of T as measured by the specificational resources $\varphi_S(T)$ must be explicitly factored in (previously, I let them simply go along for the ride). Other things being equal, universal probability bounds are better (i.e., more useful in scientific applications) the bigger they are — provided, of course that they truly are universal.

Even so, in *The Design Inference* and *No Free Lunch* I suggested that a universal probability bound is impervious any probabilistic resources that might be brought to bear against it. In those books, I offered 10^{-150} as the only probability bound anybody would ever need to draw design inferences. On the other hand, in this paper I'm saying that 10^{-120} serves that role, but that it needs to be adjusted by the specificational resources $\varphi_S(T)$, thus essentially constraining $\mathbf{P}(T|\mathbf{H})$ not by 10^{-120} but by $10^{-120} / \varphi_S(T)$. If you will, instead of a static universal probability bound of 10^{-150} , we now have a dynamic one of $10^{-120} / \varphi_S(T)$ that varies with the specificational resources $\varphi_S(T)$ and thus with the descriptive complexity of T . For many design inferences that come up in practice, it seems safe to assume that $\varphi_S(T)$ will not exceed 10^{30} (for instance, in section 7 a very generous estimate for the descriptive complexity of the bacterial flagellum came out to 10^{20}). Thus, as a rule of thumb, $10^{-120} / 10^{30} = 10^{-150}$ can still be taken as a reasonable (static) universal probability bound. At any rate, for patterns qua targets T that satisfy $\mathbf{P}(T|\mathbf{H}) \leq 10^{-150}$ and that at first blush appear to have low descriptive complexity (if only because our natural language enables us to describe them simply), the burden is on the design critic to show either that the chance hypothesis \mathbf{H} is not applicable or that $\varphi_S(T)$ is much greater than previously suspected. Getting too lucky is never a good explanation, scientific or otherwise. Thus, for practical purposes, taking 10^{-150} as a universal probability bound still works. If you will, the number stays the same, but the rationale for it has changed slightly.

One final difference that should be pointed out regarding my past work on specification is the difference between specified complexity then and now. In the past, specified complexity, as I characterized it, was a property describing a relation between a pattern and an event delineated by that pattern. Accordingly, specified complexity either obtained or did not obtain as a certain relation between a pattern and an event. In my present treatment, specified complexity still captures this relation, but it is now not merely a property but an actual number calculated by a precise formula (i.e., $\chi = -\log_2[10^{120} \cdot \varphi_S(T) \cdot \mathbf{P}(T|\mathbf{H})]$). This number can be negative, zero, or positive. When the number is greater than 1, it indicates that we are dealing with a specification.

Addendum 2: Bayesian Methods

The approach to design detection that I propose eliminates chance hypotheses when the probability of matching a suitable pattern (i.e., specification) on the basis of these chance hypotheses is small and yet an event matching the pattern still happens (i.e., the arrow lands in a small target). This eliminative approach to statistical rationality, as epitomized in Fisher's approach to significance testing, is the one most widely employed in scientific applications.⁴⁷ Nevertheless, there is an alternative approach to statistical rationality that is at odds with this eliminative approach. This is the Bayesian approach, which is essentially comparative rather than eliminative, comparing the probability of an event conditional on a chance hypothesis to its probability conditional on a design hypothesis, and preferring the hypothesis that confers the greater probability. I've argued at length elsewhere that Bayesian methods are inadequate for drawing design inferences.⁴⁸ Among the reasons I've given is the need to assess prior probabilities in employing these methods, the concomitant problem of rationally grounding these priors, and the lack of empirical grounding in estimating probabilities conditional on design hypotheses.

Here, in this addendum, I want to focus on what I regard as the most damning problem facing the Bayesian approach to design detection, namely, that it tacitly presupposes the very account of specification that it was meant to preclude. Bayesian theorists see specification as an incongruous and dispensable feature of design inferences. For instance, Timothy and Lydia McGrew, at a symposium on design reasoning, dismissed specification as having no "epistemic relevance."⁴⁹ At the same symposium, Robin Collins, also a Bayesian, remarked: "We could roughly define a specification as any type of pattern for which we have some reasons to expect an intelligent agent to produce it."⁵⁰ Thus, a Bayesian use of specification might look as follows: given some event *E* and a design hypothesis **D**, a specification would assist in inferring design for *E* if the probability of *E* conditional on **D** is increased by noting that *E* conforms to the specification.

But there's a crucial difficulty here. Consider the case of the New Jersey election commissioner Nicholas Caputo accused of rigging ballot lines. (This example appears in a number of my writings and has been widely discussed on the Internet.⁵¹ A ballot line is the order of candidates listed on a ballot. It is to the advantage of a candidate to be listed first on a ballot line because voters tend to vote more readily for such candidates.) Call Caputo's ballot line selections the event *E*. *E* consists of 41 selections of Democrats and Republicans in sequence with Democrats outnumbering Republicans 40 to 1. For definiteness, let's assume that Caputo's ballot line selections looked as follows (to my knowledge, newspapers covering the story never reported the actual sequence):

DDDDDDDDDDDDDDDDDDDDDDDDDDRDDDDDDDDDDDDDDDDDDDDDD

Thus, we suppose that for the initial 22 times, Caputo chose the Democrats to head the ballot line; then at the 23rd time, he chose the Republicans (note the underlined “R”); after which, for the remaining times, he chose the Democrats.

If Democrats and Republicans were equally likely to have come up (as Caputo claimed), this event has probability approximately 1 in 2 trillion. Improbable, yes, but by itself not enough to implicate Caputo in cheating. Highly improbable events after all happen by chance all the time — indeed, any sequence of forty-one Democrats and Republicans whatsoever would be just as unlikely (i.e., have the same “absolute specificity” in the language of section 6). What, then, do we need, additionally, to confirm cheating and thereby design? To implicate Caputo in cheating, it’s not enough merely to note a preponderance of Democrats over Republicans in some sequence of ballot line selections. Rather, one must also note that a preponderance as extreme as this is highly unlikely.

In other words, it wasn’t the event E (Caputo’s actual ballot line selections) whose improbability the Bayesian theorist needed to compute but the composite event E^* consisting of all possible ballot line selections that exhibit at least as many Democrats as Caputo selected. This event — E^* — consists of 42 possible ballot line selections and has improbability roughly 1 in 50 billion. It’s this event and this improbability on which the New Jersey Supreme Court rightly focused when it deliberated about whether Caputo had in fact cheated.⁵² Moreover, it’s this event that the Bayesian needs to identify and whose probability the Bayesian needs to compute to perform a Bayesian analysis.

But how does the Bayesian identify this event? Let’s be clear that observation never hands us composite events like E^* but only elementary outcomes like E (i.e., Caputo’s actual ballot line selection and not the ensemble of ballot line selections as extreme as Caputo’s). But whence this composite event? Within the Fisherian framework, the answer is clear: E^* is the rejection region given by counting the number of Democrats selected in 41 tries. E^* is easily seen to have high specificity in the sense of section 6 and in most contexts will exhibit specified complexity and constitute a specification.⁵³ E^* is what the court used in deciding Caputo’s case and that’s what Bayesians use. Bayesians, however, offer no account of how they identify the composite events qua specifications to which they assign probabilities. If the only events they ever considered were elementary outcomes, there would be no problem. But that’s not the case. Bayesians routinely consider such composite events. In the case of Bayesian design inferences (and Bayesians definitely want to draw a design inference with regard to Caputo’s ballot line selections), those composite events are given by specifications.

Let me paint the picture more starkly. Consider an elementary outcome E . Suppose, initially, we see no pattern that gives us reason to expect an intelligent agent produced it. But then, rummaging through our background knowledge, we suddenly discover a pattern that signifies design in E . Under a Bayesian analysis, the probability of E given the design hypothesis suddenly jumps way up. That, however, isn’t enough to allow us to infer design. As is typical in the Bayesian scheme, we need to compare a probability conditional on a design hypothesis \mathbf{D} to

one conditional on a chance hypothesis \mathbf{H} . But for which event do we compute these probabilities? As it turns out, not for the elementary outcome E , but for the composite event E^* consisting of all elementary outcomes that exhibit the pattern signifying design. In other words, we need to compare $\mathbf{P}(E^*|\mathbf{D})$ to $\mathbf{P}(E^*|\mathbf{H})$, forming the likelihood ratio $\mathbf{P}(E^*|\mathbf{D})/\mathbf{P}(E^*|\mathbf{H})$ and showing that the numerator overwhelms the denominator. The same analysis comparing $\mathbf{P}(E|\mathbf{D})$ to $\mathbf{P}(E|\mathbf{H})$ is irrelevant. Indeed, it does no good to argue for E being the result of design on the basis of some pattern unless the entire collection of elementary outcomes that exhibit that pattern (i.e., E^*) is itself improbable on the chance hypothesis.

To infer design, Bayesian methods therefore need to compare the probability of E^* conditional on the design hypothesis with the probability of E^* conditional on the chance hypothesis. E^* here is, of course, what previously we referred to as a target given by a pattern T . It follows that the Bayesian approach to statistical rationality is parasitic on the Fisherian approach and can properly adjudicate only among competing hypotheses that the Fisherian approach has thus far failed to eliminate. In particular, the Bayesian approach offers no account of how it arrives at the composite events (qua targets qua patterns qua specifications) on which it performs a Bayesian analysis. The selection of such events is highly intentional and, in the case of Bayesian design inferences, presupposes an account of specification. Specification's role in detecting design, far from being refuted by the Bayesian approach, is therefore implicit throughout Bayesian design inferences.

Endnotes

¹The three books by Plantinga on warrant are these: *Warrant: The Current Debate* (Oxford: Oxford University Press, 1993); *Warrant and Proper Function* (Oxford: Oxford University Press, 1993); *Warranted Christian Belief* (Oxford: Oxford University Press, 2000).

²Bertrand Russell, *Human Knowledge: Its Scope and Limits* (New York: Simon & Schuster, 1948), 397.

³William F. Arndt and F. Wilbur Gingrich, *A Greek-English Lexicon of the New Testament and Other Early Christian Literature* (Chicago: University of Chicago Press, 1979), 639–640.

⁴William A. Dembski, *Design Inference: Eliminating Chance through Small Probabilities* (Cambridge: Cambridge University Press, 1998) and *No Free Lunch: Why Specified Complexity Cannot Be Purchased Without Intelligence* (Lanham, Md.: Rowman and Littlefield, 2002).

⁵Charles Peirce, “The General Theory of Probable Inference,” in Justus Buchler, ed., *Philosophical Writings of Peirce*, 190–217 (1883; reprinted New York: Dover, 1955), 206–211.

⁶For a brief summary of Fisher’s views on tests of significance and null hypotheses, see Ronald A. Fisher, *The Design of Experiments* (New York: Hafner, 1935), 13–17.

⁷Colin Howson and Peter Urbach, *Scientific Reasoning: The Bayesian Approach*, 2nd ed. (LaSalle, Ill.: Open Court, 1993), 178. For a similar criticism see Ian Hacking, *Logic of Statistical Inference* (Cambridge: Cambridge University Press, 1965), 81–83.

⁸Dembski, *Design Inference*, ch. 6.

⁹A false positive here is the error of eliminating a chance hypothesis as the explanation of an event when the chance hypothesis actually is in force. Statisticians refer to such false positives as “type I errors.”

¹⁰See Dembski, *The Design Inference*, 199–203.

¹¹For a full account of such probabilities, see my article “Uniform Probability,” *Journal of Theoretical Probability* 3(4) (1990): 611–626. On the real line, this measure is just Lebesgue measure.

¹²Strictly speaking, $\mathbf{P}(\cdot|\mathbf{H})$ can be represented as $f d\mathbf{U}$ only if $\mathbf{P}(\cdot|\mathbf{H})$ is absolutely continuous with respect to \mathbf{U} , i.e., all the subsets of Ω that have zero probability with respect to \mathbf{U} must also have zero probability with respect to $\mathbf{P}(\cdot|\mathbf{H})$ (this is just the Radon-Nikodym Theorem — see Heinz Bauer, *Probability Theory and Elements of Measure Theory*, trans. R. B. Burckel, 2nd English ed. [New York: Academic Press, 1981], 86). Nonetheless, when Ω is finite, any probability measure whatsoever is absolutely continuous with respect to the uniform probability. What’s more, in most cases where Ω is infinite, it is possible to approximate probability measures of the form $\mathbf{P}(\cdot|\mathbf{H})$ arbitrarily closely by probability measures of the form $f d\mathbf{U}$ (the latter probabilities can be made to “converge weakly” to the former — see Patrick Billingsley, *Convergence of Probability Measures*, 2nd ed. [New York: Wiley, 1999], ch. 1). Physicists and engineers bypass questions of absolute continuity and weak convergence by thinking of f as a generalized function — see, for instance, Paul Dirac, *The Principles of Quantum Mechanics*, 4th ed. (Oxford: Oxford University Press, 1958), 48 or Michael Reed and Barry Simon, *Methods of Modern Mathematical Physics I: Functional Analysis*, revised and enlarged (New York: Academic Press, 1980), 148.

¹³William Feller, *An Introduction to Probability Theory and Its Applications*, vol. 1, 3rd ed. (New York: Wiley, 1968), 167–169.

¹⁴The statistics text in question is David Freedman, Robert Pisani, and Roger Purves, *Statistics* (New York: Norton, 1978), 426–427. Fisher’s original account can be found in Ronald A. Fisher, *Experiments in Plant Hybridisation* (Edinburgh: Oliver and Boyd, 1965), 53. For a more recent reevaluation of Mendel’s data, which still concludes that “the segregations are in general closer to Mendel’s expectations than chance would dictate,” see A. W. F. Edwards, “Are Mendel’s Results Really Too Close?” *Biological Review* 61 (1986): 295–312.

¹⁵See Bauer, *Probability Theory*, 44. Indicator functions are also known as characteristic functions.

¹⁶Gregory J. Chaitin, "On the Length of Programs for Computing Finite Binary Sequences," *Journal of the Association for Computing Machinery* 13 (1966): 547–569; Andrei Kolmogorov, "Three Approaches to the Quantitative Definition of Information," *Problemy Peredachi Informatsii* (in translation) 1(1) (1965): 3–11; Ray J. Solomonoff, "A Formal Theory of Inductive Inference, Part I," *Information and Control* 7 (1964): 1–22 and Ray J. Solomonoff, "A Formal Theory of Inductive Inference, Part II," *Information and Control* 7 (1964): 224–254.

¹⁷See Hartley Rogers Jr., *Theory of Recursive Functions and Effective Computability* (1967; reprinted Cambridge, Mass.: MIT Press, 1987).

¹⁸For an overview of the algorithmic information theoretic approach to randomness, see Peter Smith, *Explaining Chaos* (Cambridge: Cambridge University Press, 1998), ch. 9.

¹⁹See C. H. Woo, "Laws and Boundary Conditions," in *Complexity, Entropy and the Physics of Information*, ed. W. H. Zurek (Reading, Mass.: Addison-Wesley, 1990), 132–133, where Woo offers some preliminary remarks about the connection between thermodynamic and algorithmic entropy.

²⁰Christian de Duve offers such arguments in *Vital Dust: Life as a Cosmic Imperative* (New York: Basic Books, 1995).

²¹The proof is straightforward: In 100 coin tosses, on average half will repeat the previous toss, implying about 50 two-repetitions. Of these 50 two-repetitions, on average half will repeat the previous toss, implying about 25 three-repetitions. Continuing in this vein, we find on average 12 four-repetitions, 6 five-repetitions, 3 six-repetitions, and 1 seven-repetition. See Ivars Peterson, *The Jungles of Randomness: A Mathematical Safari* (New York: Wiley, 1998), 5.

²²G. H. Hardy and E. M. Wright, *An Introduction to the Theory of Numbers*, 5th ed. (Oxford: Clarendon Press, 1979), 128.

²³It follows that specification is intimately connected with discussions in the self-organizational literature about the "edge of chaos," in which interesting self-organizational events happen not where things are completely chaotic (i.e., entirely chance-driven and thus not easily describable) and not where things are completely ordered (i.e., so predictable as to be easily describable but therefore not improbable). See Roger Lewin, *Complexity: Life at the Edge of Chaos*, 2nd ed. (Chicago: University of Chicago Press, 2000).

²⁴There is a well-established theory of descriptive complexity, which takes as its point of departure Chaitin-Kolmogorov-Solomonoff theory of bit-string compressibility, namely, the theory of Minimal Description Length (MDL). The fundamental idea behind MDL is that order in data "can be used to *compress* the data, i.e., to describe it using fewer symbols than needed to describe the data literally." See <http://www.mdl-research.org> (last accessed June 17, 2005).

²⁵For an account of such complexity measures, see Dembski, *The Design Inference*, ch. 4.

²⁶In characterizing how we eliminate chance and infer design, I'm providing what philosophers call a "rational reconstruction." That is, I'm providing a theoretical framework for something we humans do without conscious attention to theoretical or formal principles. As I see it, semiotic agents like us tacitly assess the descriptive complexity of patterns on the basis of their own personal background knowledge. It is an interesting question, however, whether the complexity of patterns need not be relativized to such agents. Robert Koons considers this possibility by developing a notion of *ontological complexity*. See Robert Koons, "Are Probabilities Indispensable to the Design Inference," *Progress in Complexity, Information, and Design* 1(1) (2002): available online at http://www.iscid.org/papers/Koons_AreProbabilities_112701.pdf (last accessed June 21, 2005).

²⁷See Dembski, *No Free Lunch*, ch. 3.

²⁸See Josh English's webpage titled "The Odds of Poker," <http://www.spiritone.com/~english/words/poker.html> (last accessed June 18, 2005).

²⁹Dan Brown, *The Da Vinci Code* (New York: Doubleday, 2003), 187–189.

³⁰For the mathematics of the Fibonacci sequence, see G. H. Hardy and E. M. Wright, *An Introduction to the Theory of Numbers*, 5th ed. (Oxford: Clarendon Press, 1979), 148–153. For an application of this sequence to biology, see Ian Stewart, *Life's Other Secret: The New Mathematics of the Living World* (New York: Wiley, 1998), 122–132.

³¹Seth Lloyd, "Computational Capacity of the Universe," *Physical Review Letters* 88(23) (2002): 7901–4.

³²William A. Dembski, "The Chance of the Gaps," in Neil Manson, ed., *God and Design: The Teleological Argument and Modern Science* (London: Routledge, 2002), 251–274.

³³See my article "Irreducible Complexity Revisited" at www.designinference.com (last accessed June 17, 2005). See also section 5.10 of my book *No Free Lunch*.

³⁴Abraham de Moivre, *The Doctrine of Chances* (1718; reprinted New York: Chelsea, 1967), v.

³⁵See William A. Dembski, *The Design Revolution: Answering the Toughest Questions about Intelligent Design* (Downers Grove, Ill.: InterVarsity, 2004), ch. 33.

³⁶See Jonathan Wells, *Icons of Evolution* (Washington, DC: Regnery, 2000), ch. 2.

³⁷Consider the following commentary on this fact by Darwinist Michael Ruse: "At the moment the hand of human design and intention hangs heavily over everything [i.e., everything in origin-of-life research], but work is going forward rapidly to create conditions in which molecules can make the right and needed steps without constant outside help. When that happens, as one researcher has put it, 'the dreaming stops and the fun begins,' as one looks to see if such a replicating RNA molecule could make a DNA molecule, and how it might function both with respect to producing new cell components and with respect to transmitting information from one generation to the next." Quoted in *Can A Darwinian Be a Christian?* (Cambridge: Cambridge University Press, 2001), 64. The point I wish to underscore is that, to date, the dreaming continues and shows no signs of stopping.

³⁸Showing that the Darwinian mechanism renders chance the winner over design in biological evolution is the point of Richard Dawkins's *Climbing Mount Improbable* (New York: Norton, 1996).

³⁹Gary Jason, *Critical Thinking: Developing an Effective Worldview* (Belmont, Calif.: Wadsworth, 2001), 133.

⁴⁰John Earman, *Bayes or Bust? A Critical Examination of Bayesian Confirmation Theory* (Cambridge, Mass.: MIT Press, 1992), ch. 7, titled "A Plea for Eliminative Induction."

⁴¹See William A. Dembski, "The Logical Underpinnings of Intelligent Design," in W. A. Dembski and M. Ruse, eds., *Debating Design: From Darwin to DNA*, 311–330 (Cambridge: Cambridge University Press, 2004), 328–329.

⁴²For a fuller treatment of this line of criticism as well as a response to it, see Dembski, *The Design Revolution*, ch. 26.

⁴³Elliott Sober, "Testability," *Proceedings and Addresses of the American Philosophical Association* 73(2) (1999): 73, n. 20. For the record, Sober embraces the independent knowledge requirement.

⁴⁴For Sober's defense of the independent knowledge requirement as well as my response to it, see my book *The Design Revolution*, ch. 32.

⁴⁵Rob Koons was the first to raise the possibility that specification might only require tractability. He did this at the Symposium on Design Reasoning, Calvin College, May 2001. At the time, I was attracted to this possibility, but couldn't see my way clear to it because I was still committed to seeing specifications as incorporating

prespecifications and thus requiring tractability to apply, in an artificial way, to prespecifications (the artificiality being that prespecifications would satisfy the conditional independence condition simply by being identified prior to the event in question and then satisfy the tractability condition trivially by providing information about themselves). See Dembski, *The Design Inference*, sec. 5.3.

⁴⁶My previous universal probability bound of 10^{-150} was calculated using considerations similar to those of Seth Lloyd: it looked to the number of state changes in elementary particles throughout the history of the known, observable universe, a number I calculated at 10^{150} . Lloyd's approach is more elegant and employs deeper insights into physics. In consequence, his approach yields a more precise estimate for the universal probability bound.

⁴⁷Even the critics of eliminative approaches admit as much. See Colin Howson and Peter Urbach, *Scientific Reasoning: The Bayesian Approach*, 2nd ed. (LaSalle, Ill.: Open Court, 1993), 192 and Richard M. Royall, *Statistical Evidence: A Likelihood Paradigm* (London: Chapman & Hall, 1997), 61–62.

⁴⁸See Dembski, *The Design Revolution*, ch. 33 as well as my article “Detecting Design by Eliminating Chance: A Response to Robin Collins,” *Christian Scholar's Review* 30(3) (Spring 2001): 343–357.

⁴⁹”Symposium on Design Reasoning, Calvin College, May 2001. Transcripts of this symposium will at some point be appearing on the International Society for Complexity, Information, and Design's website (www.iscid.org).

⁵⁰Ibid.

⁵¹See, for instance, Dembski, *No Free Lunch*, 79–83.

⁵²For the *New York Times* account that makes this clear, see the 23 July 1985 issue, page B1. For a statistics textbook that analyzes the Caputo case, see David S. Moore and George P. McCabe, *Introduction to the Practice of Statistics*, 2nd ed. (New York: W. H. Freeman, 1993), 376–377.

⁵³For the details, see Dembski, *No Free Lunch*, 79–83.